CSELAB

Computation Science and Engineering Laboratory
www.cse-lab.ch

ETH zürich

Uncertainty Quantification in Cloud Cavitation Collapse using Multi-Level Monte Carlo

Jonas Šukys

Computational Science and Engineering Laboratory
Swiss Federal Aquatic Research Laboratory
ETH Domain, Zurich, Switzerland

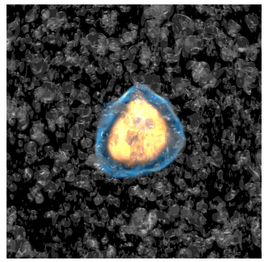

Advanced Modeling & Simulation Seminar Series
NASA Ames Research Center, California, United States

January 12, 2017

Work in progress in collaboration with

Petros Koumoutsakos

Ursula Rasthofer

Panagiotis Hadjidoukas

Diego Rossinelli

Fabian Wermelinger

Cavitation phenomenon

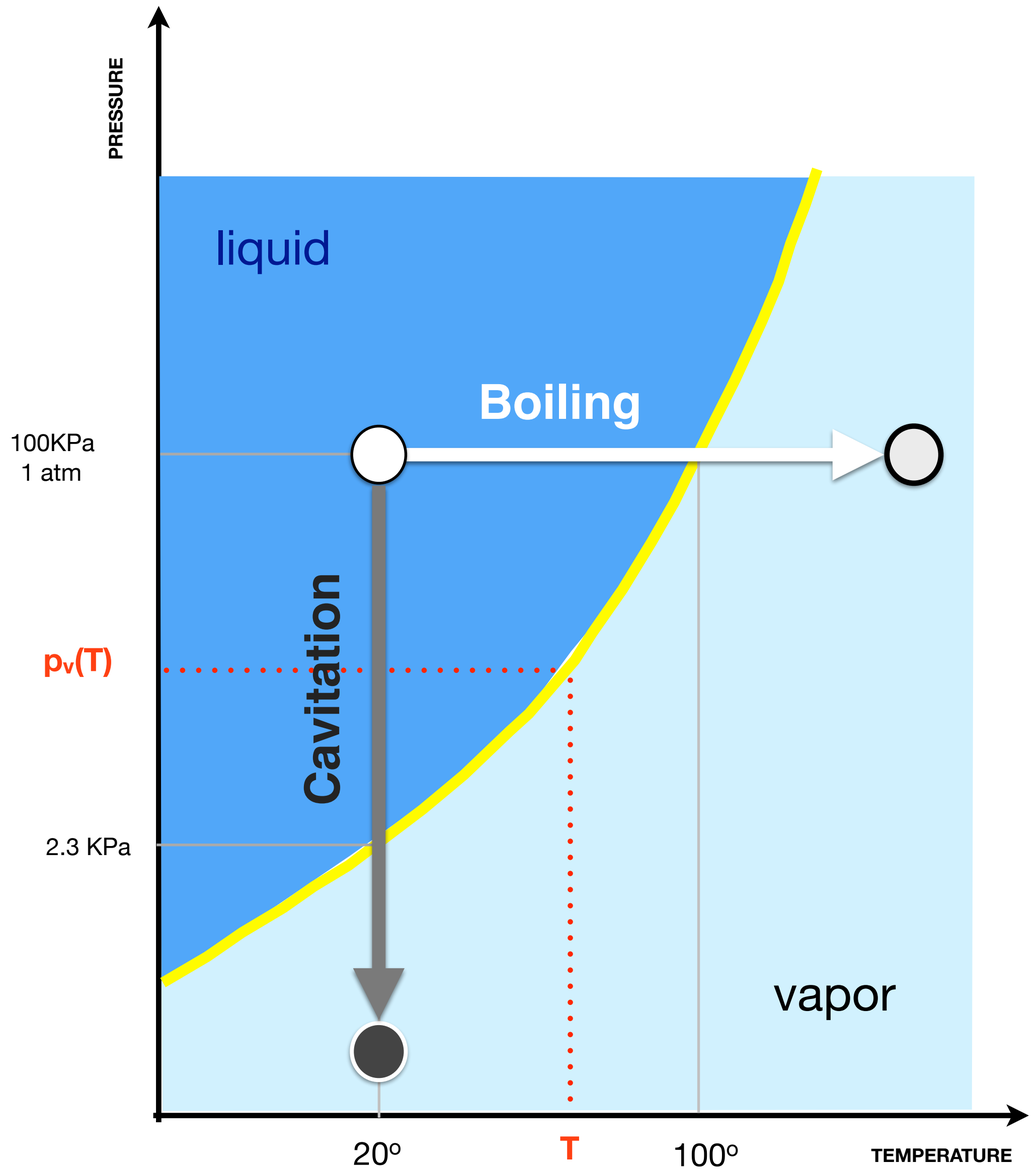

Image courtesy:
C. Koumoutsakos

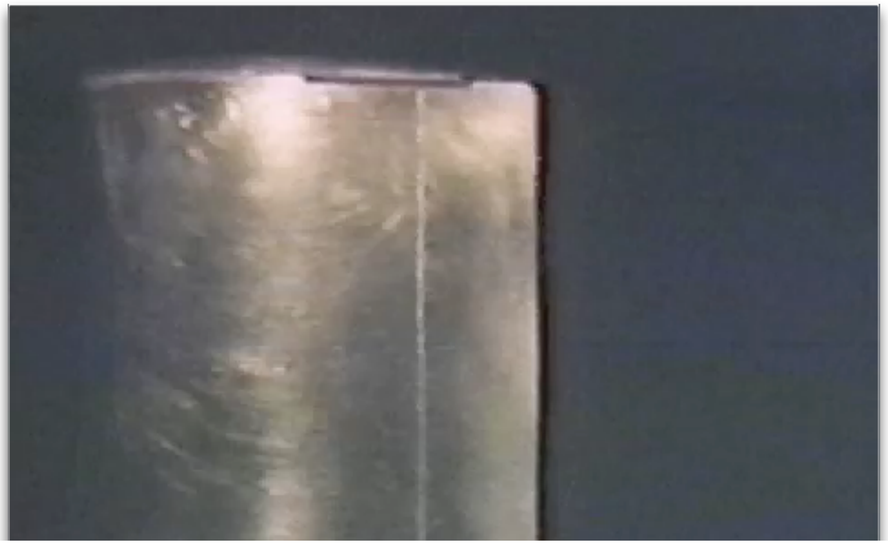

$$p + \frac{1}{2}\rho u^2 = \text{const.}$$

Image courtesy:
C. Brennen

Single cavity bubble collapse

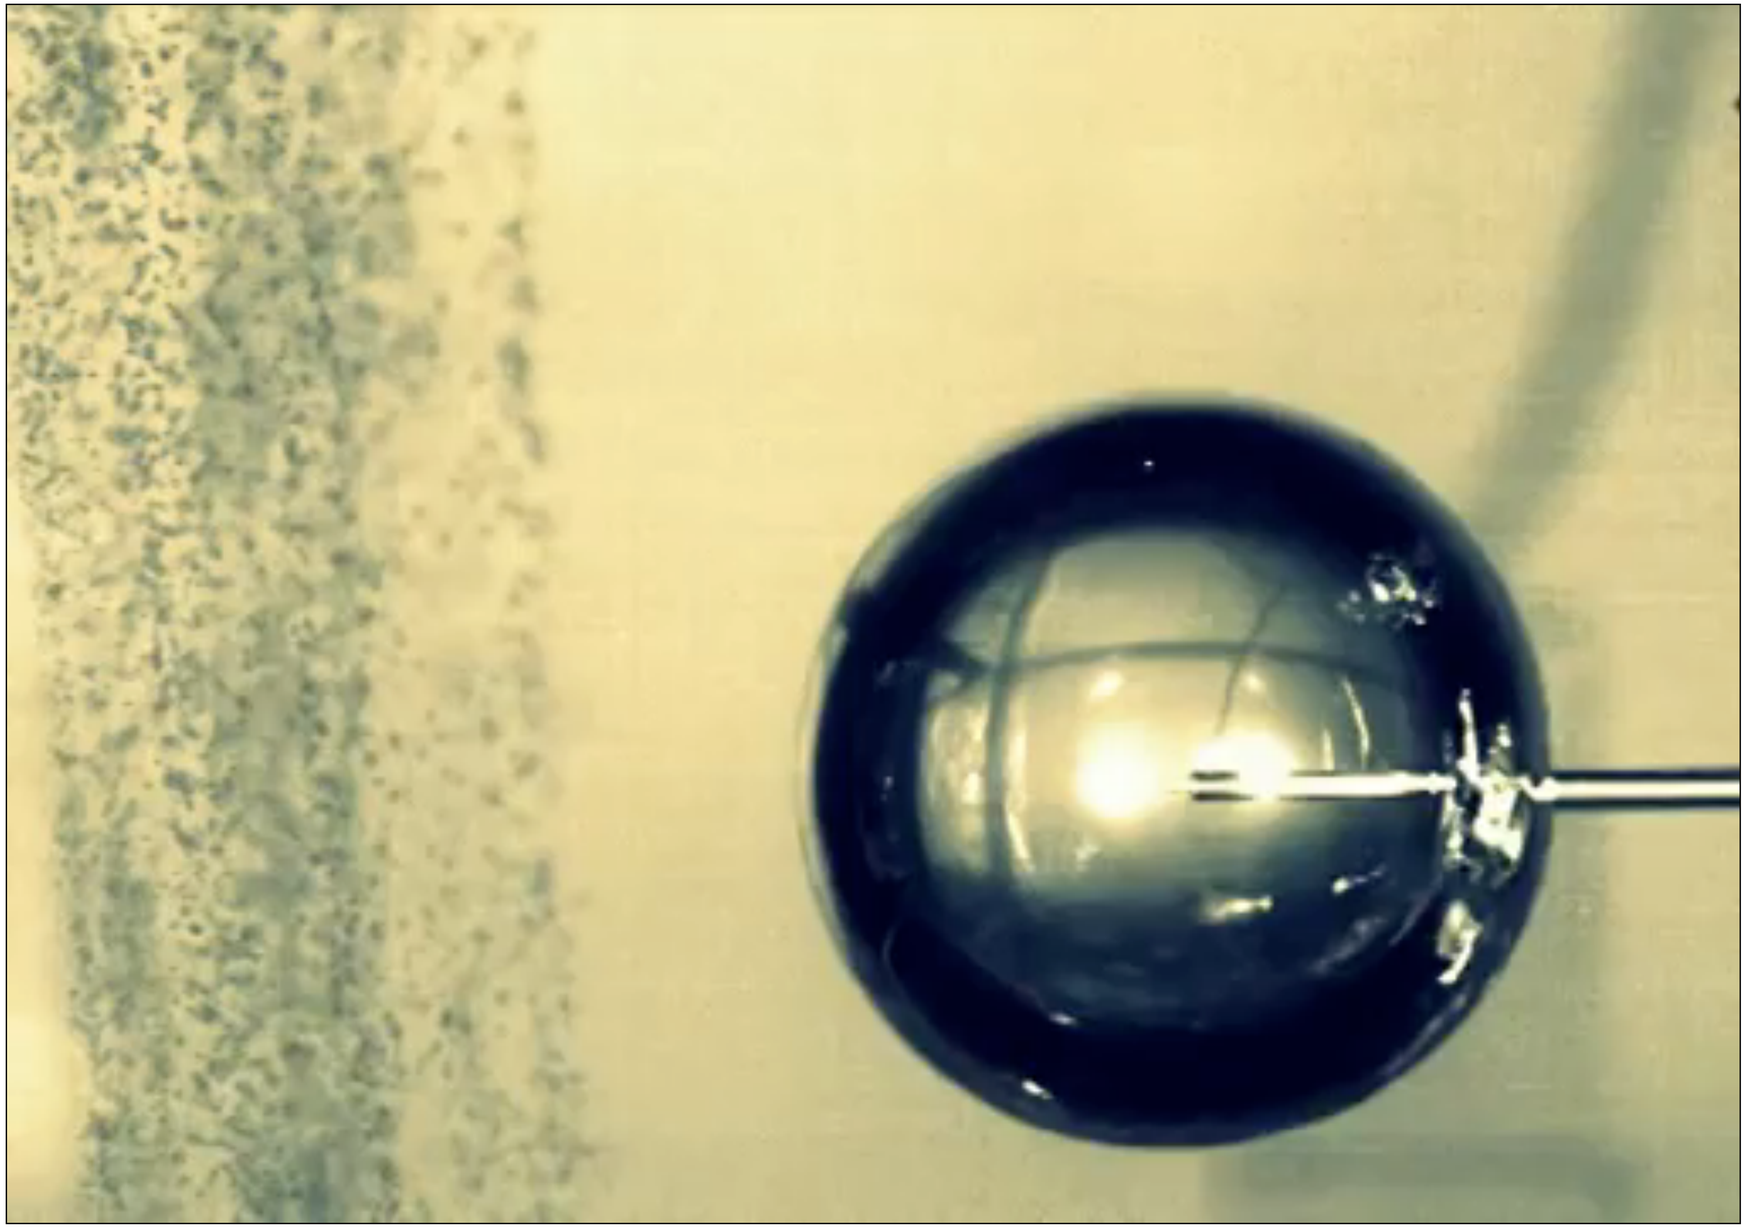

Image courtesy: DynaFlow Inc.

Destructive power of cavitation

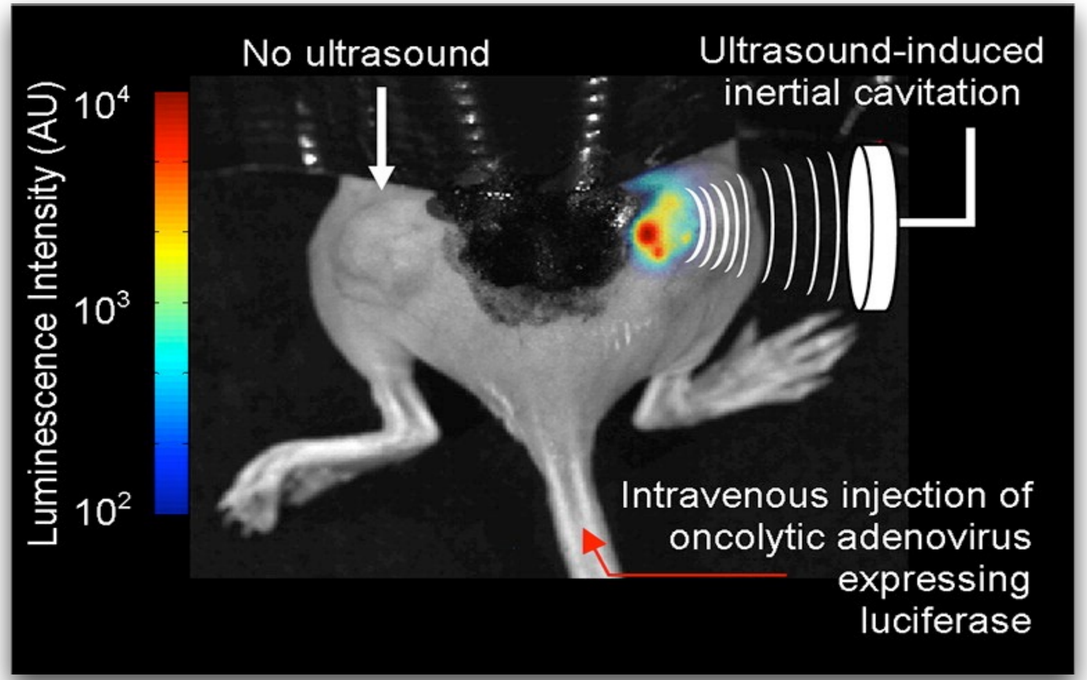

AVOID to maintain performance

- turbines (hydroelectricity, pumps)
- high pressure fuel injectors
- high pressure pipes
- propellers

Image courtesy:
Brennen, "Hydrodynamics of Pumps". Oxford University Press, 1994.

HARNESS for medical treatments

- ultrasonic drug delivery
- kidney shockwave lithotripsy
  - collapse of cavities near stone surface

Image courtesy:
Bazan-Peregrino et al., Cavitation-enhanced delivery of a replicating oncolytic adenovirus to tumors using focused ultrasound.
Journal of Controlled Release Volume 169, Issues 1–2, 2013, pp. 40 - 47.

Prevalent configurations of cavity **clouds**

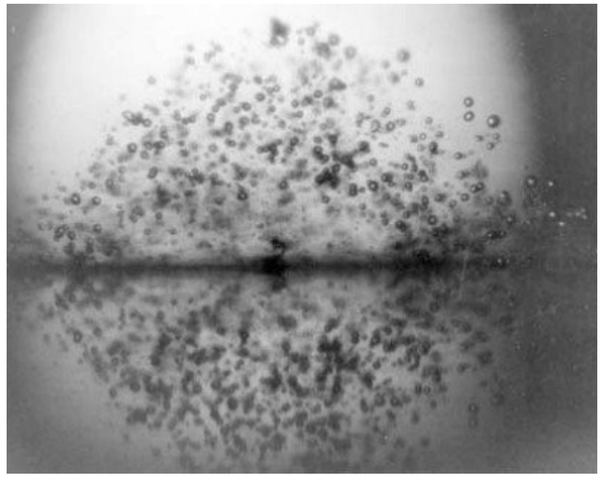

Image courtesy: Plesset and Ellis (1955)

free field

Image courtesy: Brennen (1970)

past obstacles

en.wikipedia.org/wiki/Rayleigh-Plesset_equation

propeller blades
turbine blades

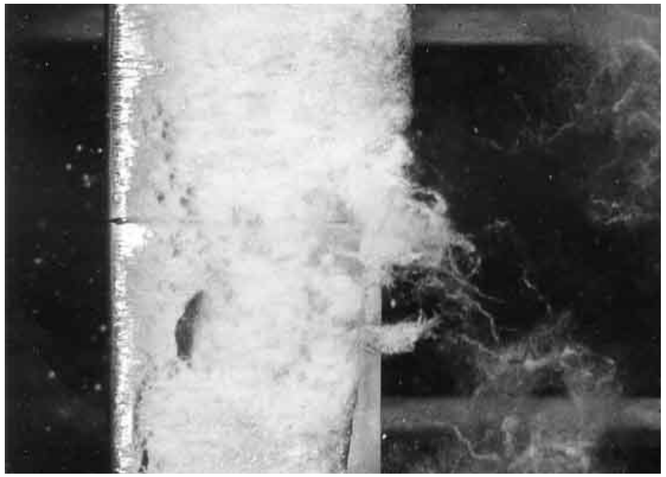

Image courtesy of Kermeen (1956)

above hydrofoils

- **clouds** of many (thousands) cavity bubbles
- **interactions** of the bubbles play a key role

State of the art

EXPERIMENTS

- Cloud interaction parameter, collapse time to radius (Brennen et al.)
- Averaged quantities, damage assessments (Lohse, Keller, Bose et al.)
- **Single/double** bubble, proximity effects on jetting (Tomita and Shima)

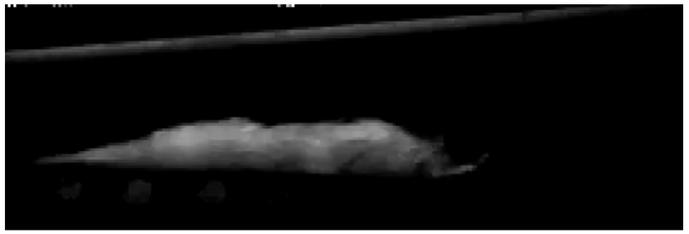

THEORY/MODELS

- **Single** bubble, radial symmetry (ODE):
  Rayleigh-Plesset (1949), Prosperetti-Lezzi (1985), Geers (2014)

SIMULATIONS

- **Single** bubble (Colonius, Caltech), multiple bubbles with models
- **Clouds** 120 bubbles, under-resolving and coarse-graining (Adams, TUM)
- **Clouds** 80 bubbles, k-div terms, interface sharpening (Tiwari, et al., 2015)
- **Clouds** 10K bubbles - Rosinelli, Hejazialhosseini, Hadjidoukas, et al. (2013)
- **Clouds** 50K bubbles - Šukys, Hadjidoukas, Rasthofer, Wermelinger, et al. (2016)

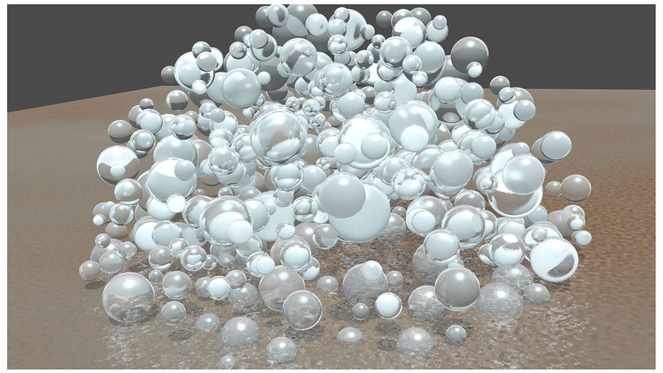

UNCERTAINTY QUANTIFICATION IN CAVITATION

- Congedo, Goncalves, Rodio (2015)
- 2D, sDEM [Abgrall, 2015], forward UQ propagation

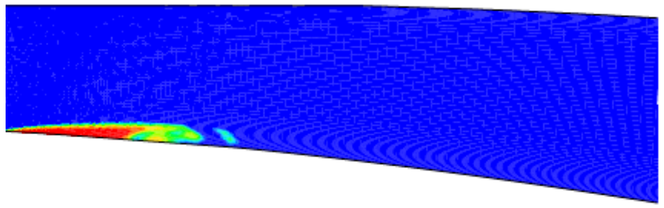

Governing equations [Kappila] [Masoni] [Allaire]

Multiphase flow equations

$$\begin{cases} (\alpha_1\rho_1)_t + \nabla\cdot(\alpha_1\rho_1\mathbf{u}) = 0, \\ (\alpha_2\rho_2)_t + \nabla\cdot(\alpha_2\rho_2\mathbf{u}) = 0, \\ (\rho\mathbf{u})_t + \nabla\cdot(\rho\mathbf{u}\otimes\mathbf{u} + p\mathbf{I}) = 0, \\ E_t + \nabla\cdot((E+p)\mathbf{u}) = 0. \end{cases}$$

Advection of phase volume fractions

$$(\alpha_2)_t + \mathbf{u}\cdot\nabla\alpha_2 = K(\alpha_{1,2}, \rho_{1,2}, c_{1,2})\nabla\cdot\mathbf{u}.$$

Equation of state (water phase: stiffened)

$$E = \frac{1}{2}\rho\mathbf{u}^2 + \Gamma p + \Pi, \quad \Gamma = \frac{1}{\gamma-1}, \quad \Pi = \frac{\gamma p_c}{\gamma-1}.$$

$$p = \frac{(E-\rho\mathbf{u}^2) - (\alpha_1\Pi_1 + \alpha_2\Pi_2)}{\alpha_1\Gamma_1 + \alpha_2\Gamma_2}, \quad \frac{1}{\rho c^2} = \frac{\alpha_1}{\rho_1 c_1^2} + \frac{\alpha_2}{\rho_2 c_2^2}.$$

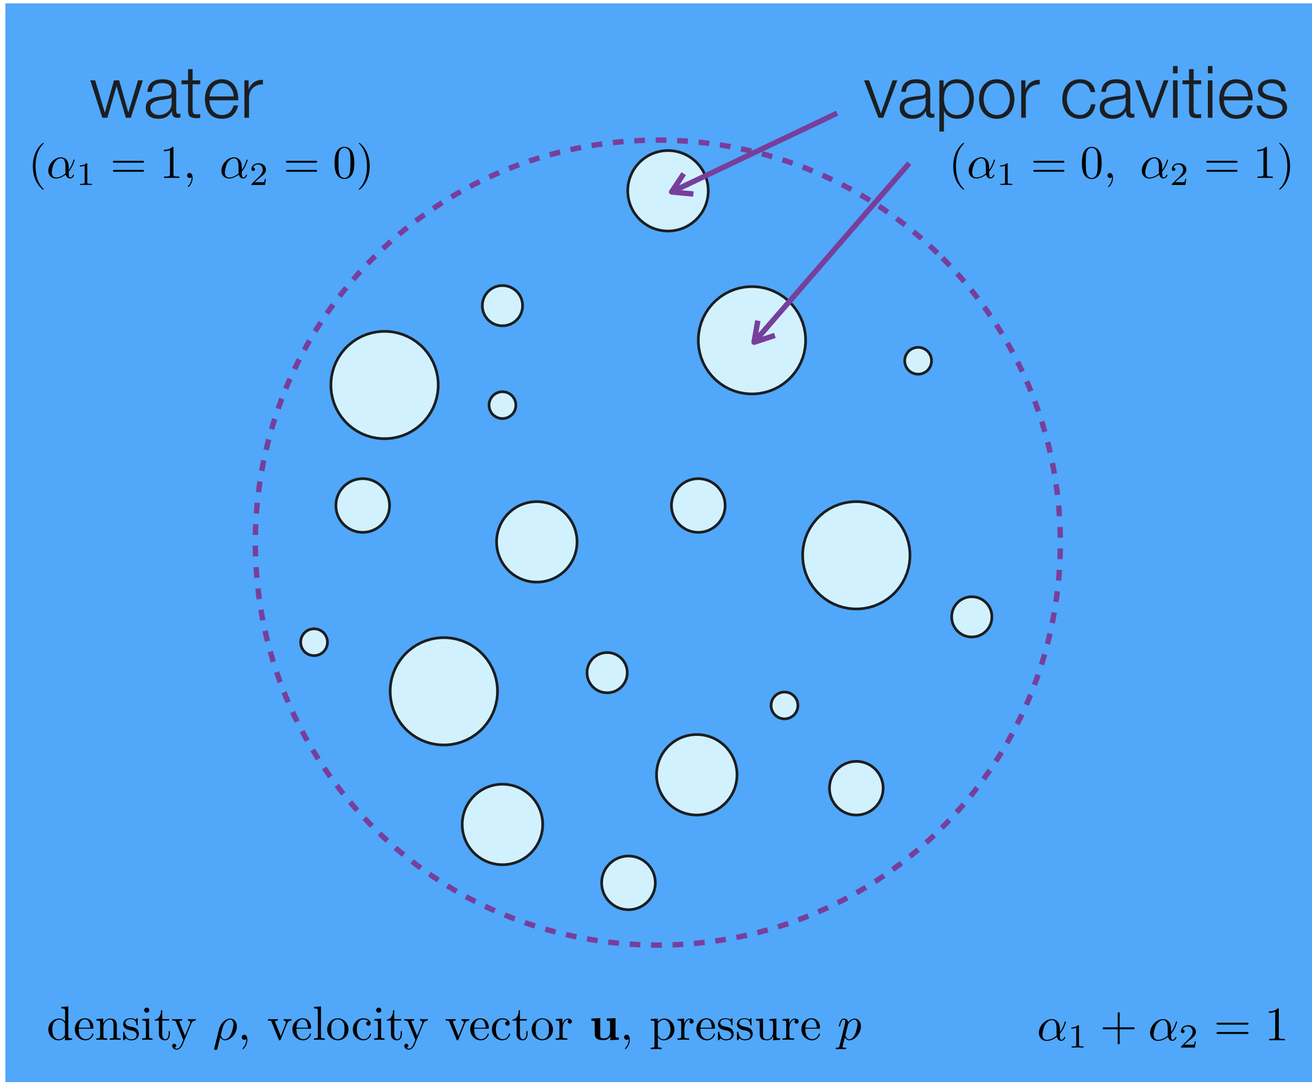

2D slice of a 3D domain

cavity sizes of 50-200 μm (log-Gaussian)

Problem setup

Generation of the cavity cloud

- **locations:** uniform distribution
- **radii:** log-Gaussian, 50 - 200 μm

	water	vapor
$\rho\ [kg/m^3]$	1000	1
$p\ [bar]$	100	0.0234
γ	6.59	1.4
$p_c\ [bar]$	4049	0

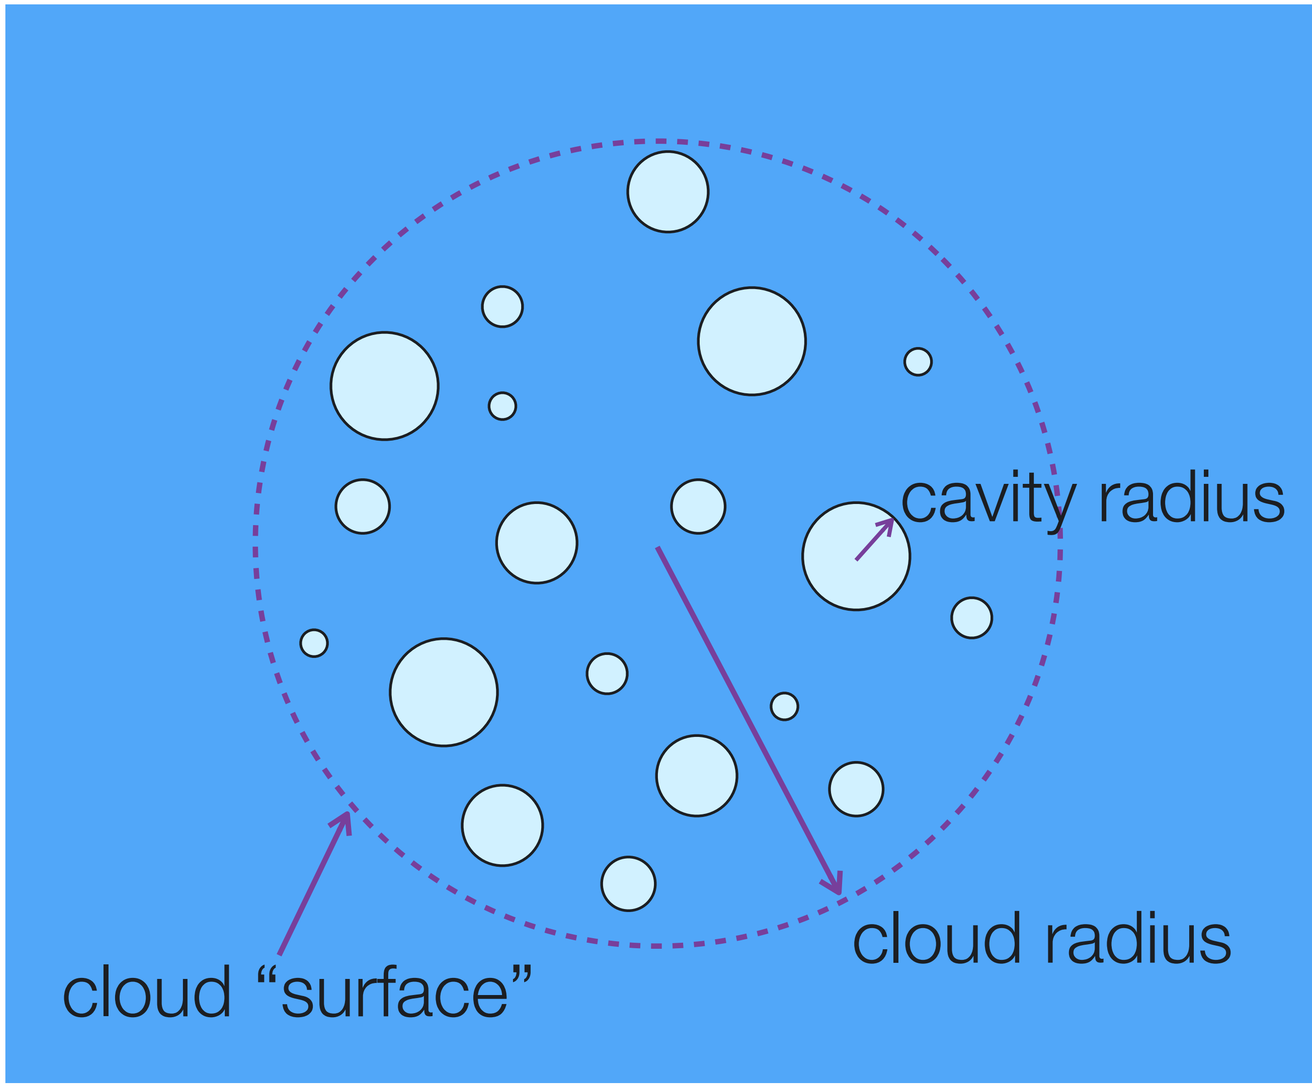

Finite Volume Solver

$$\partial_t \mathbf{U}(\mathbf{x},t) + \operatorname{div} \mathbf{F}(\mathbf{U},\mathbf{x}) = 0$$

- Cell averages

$$\mathbf{U}_j(t) \approx \frac{1}{|C_j|} \int_{C_j} \mathbf{U}(x,t)dx$$

- Semi-discrete formulation (ODE)

$$\frac{\mathrm{d}}{\mathrm{dt}} \mathbf{U}_j(t) + \frac{1}{\Delta x} \left(\mathbf{F}_{j+\frac{1}{2}} - \mathbf{F}_{j-\frac{1}{2}} \right) = 0$$

- High order reconstruction

WENO3 / WENO5
[Harten, Shu, Osher]

- Approximate Riemann solver HLLC

$$\mathbf{F}_{j+\frac{1}{2}} \approx \mathbf{F}^{\mathrm{HLLC}}_{j+\frac{1}{2}}(\mathbf{U}^+, \mathbf{U}^-)$$

- RK3 time stepping [Gottlieb, Shu, Tadmor]

$$\mathbf{U}^n_j \rightarrow \mathbf{U}^{n+1}_j$$

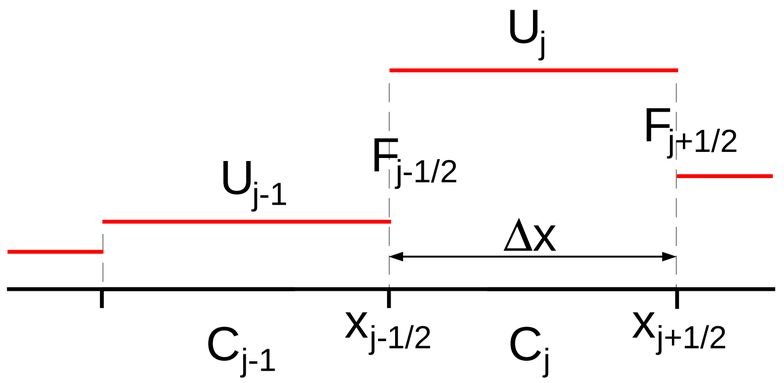

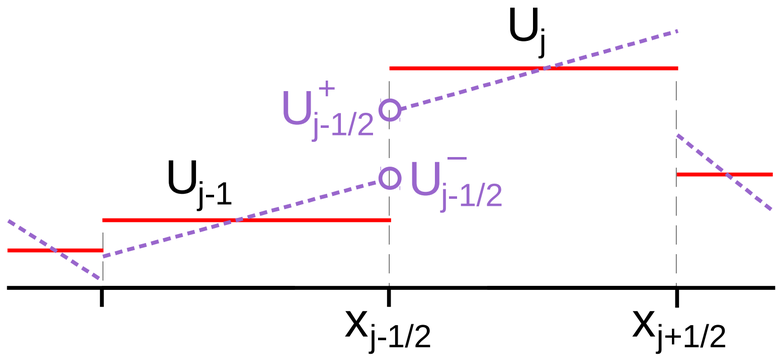

CUBISM-MPCF

Peta-scale Multi-Phase Compressible Flow approximate Riemann solver

[Rossinelli, Hejazialhosseini, Hadjidoukas, Conti, Bergdorf, Wermelinger, Rasthofer, Šukys]

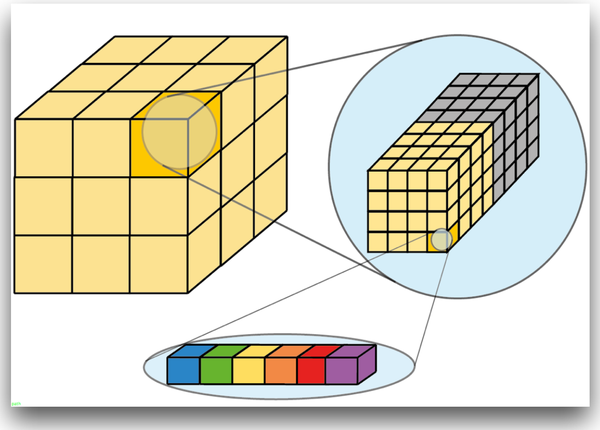

Block-based memory layout

(spatial locality)

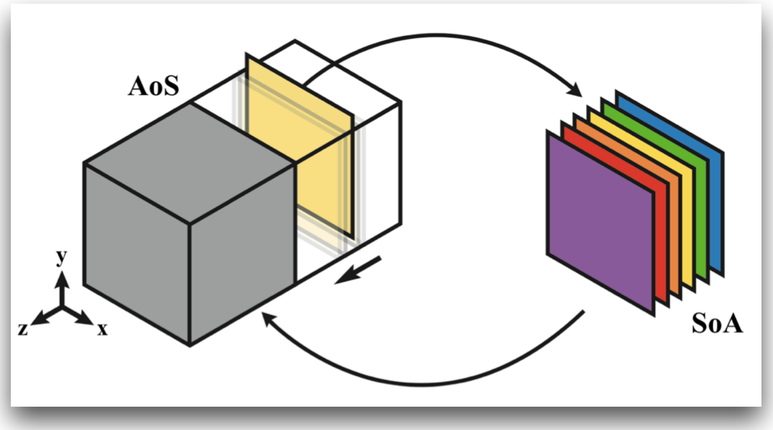

Instruction/data-level parallelism

(**S**tructure **o**f **A**rrays for SSE/QPX vectorization)

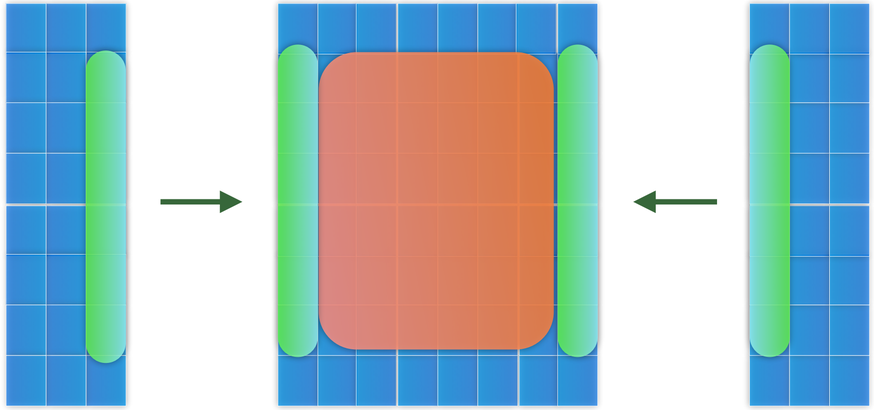

Domain decomposition MPI/OpenMP

(dynamic loop scheduling) (non-blocking P2P communication)

(asynchronous progress for C/T overlap)

- ACM Gordon Bell Prize: **14.4 Pflops (72% peak)** on Sequoia (IBM BlueGene/Q, 1.6M cores)
- Wavelet-based I/O **compression** | ~100x reduction | 1% overhead
- **Fault-tolerance with restart** mechanism | lossless compression ~10x reduction

Petascale simulations of cloud cavitation collapse

50 thousand cavities at 100 bar
0.5 billion mesh elements
25 thousand time steps
25x pressure amplification

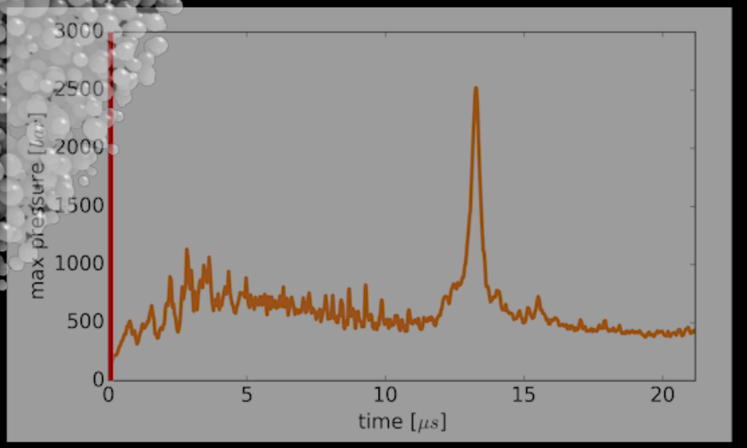

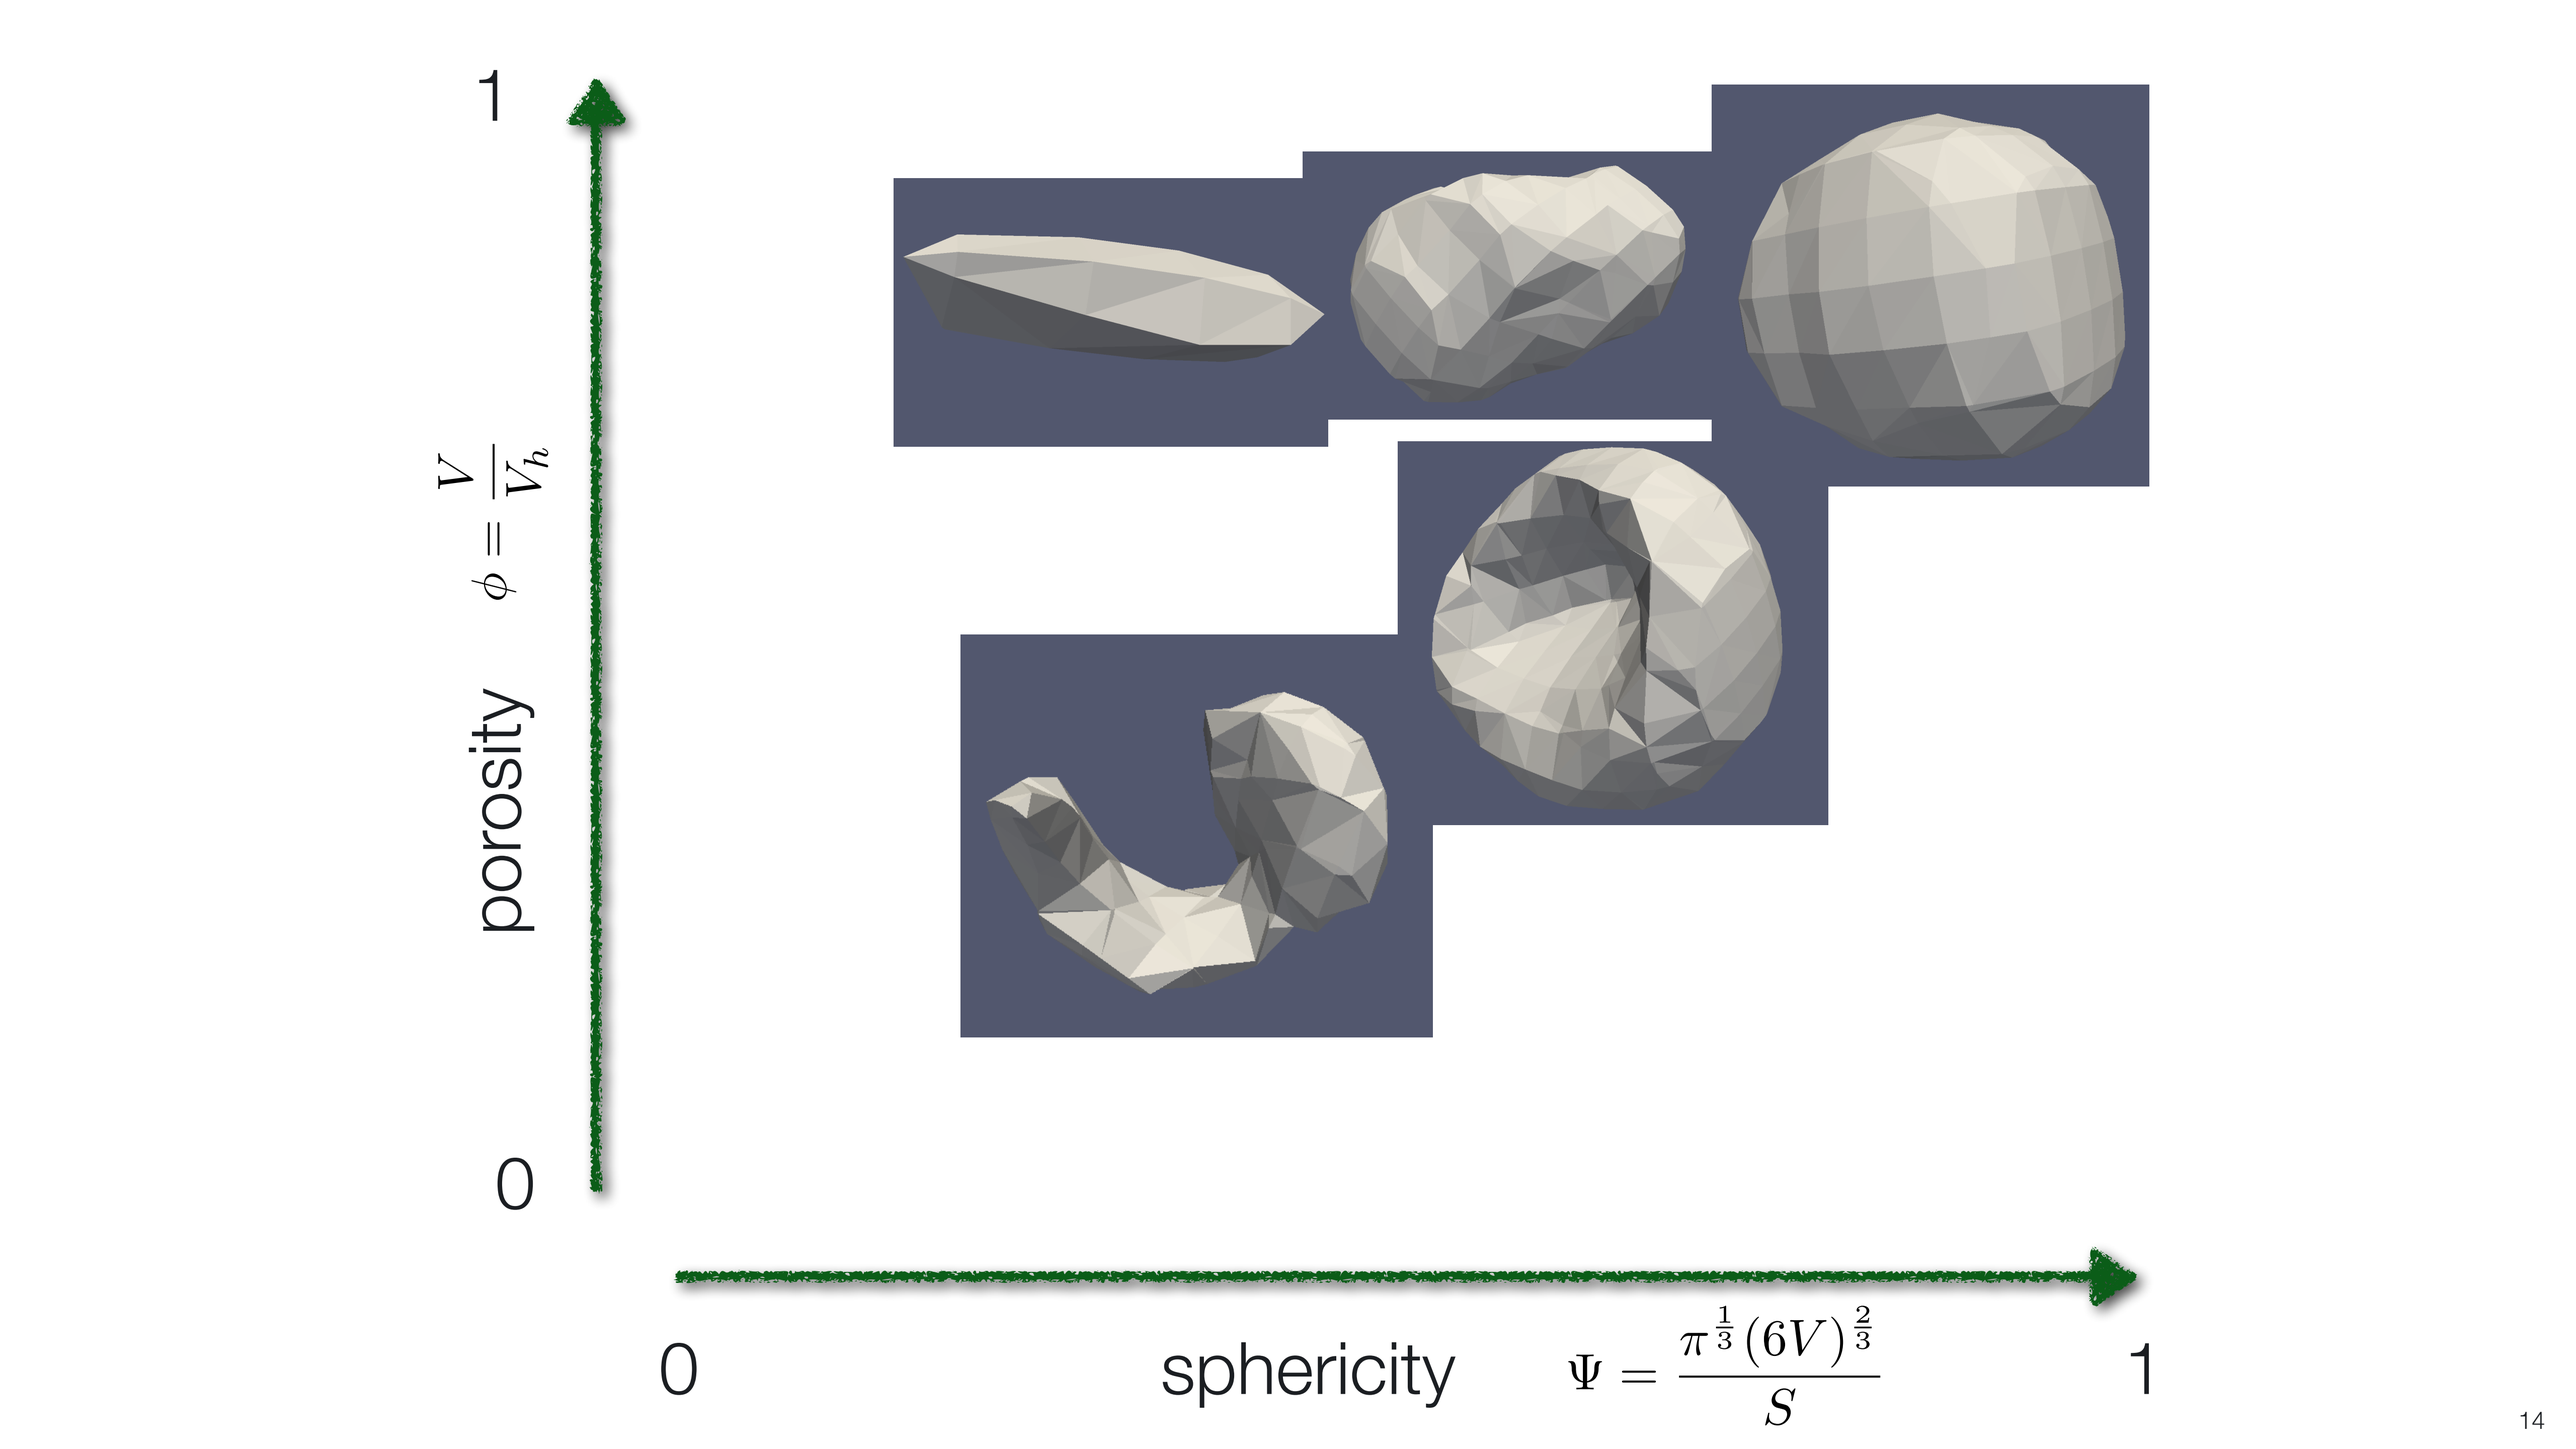
1
porosity $\phi = \frac{V}{V_h}$
0
0
sphericity $\Psi = \frac{\pi^{\frac{1}{3}}(6V)^{\frac{2}{3}}}{S}$
1

Averages | non-sphericities and non-porosities

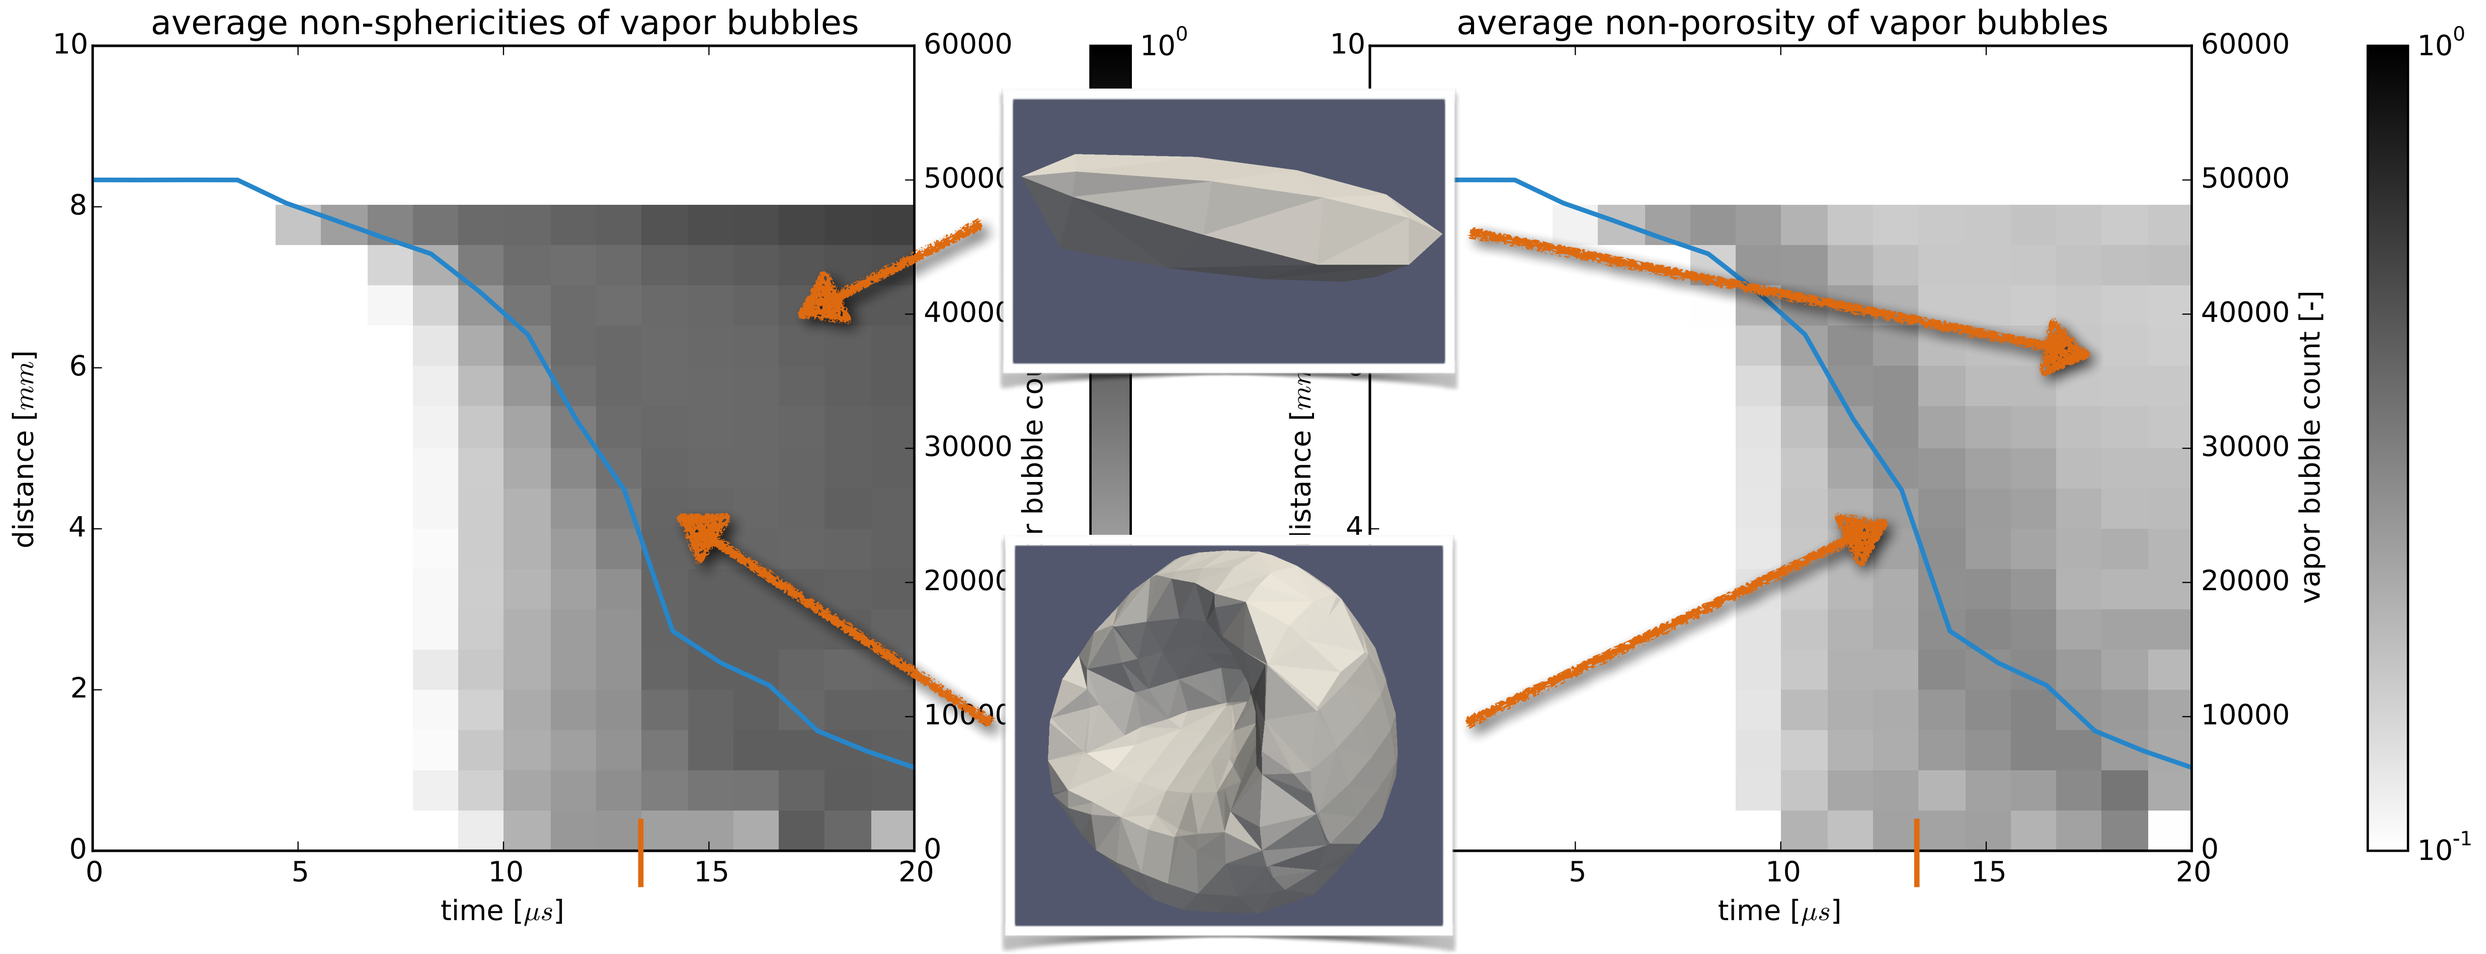

propagation of non-spherical cavities from cloud surface to center

propagation of non-porous cavities from cloud surface to center

Holes in the cavities from re-entrant micro-jets

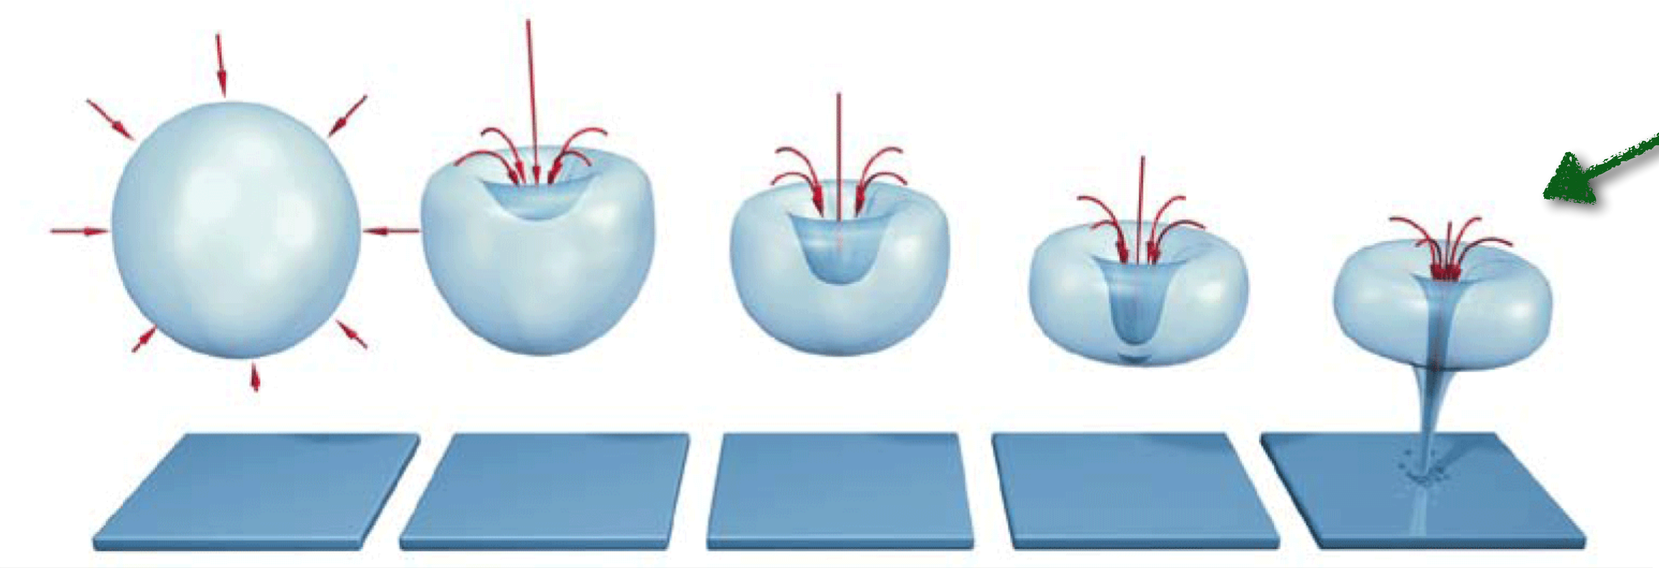

caused by re-entrant micro-jet

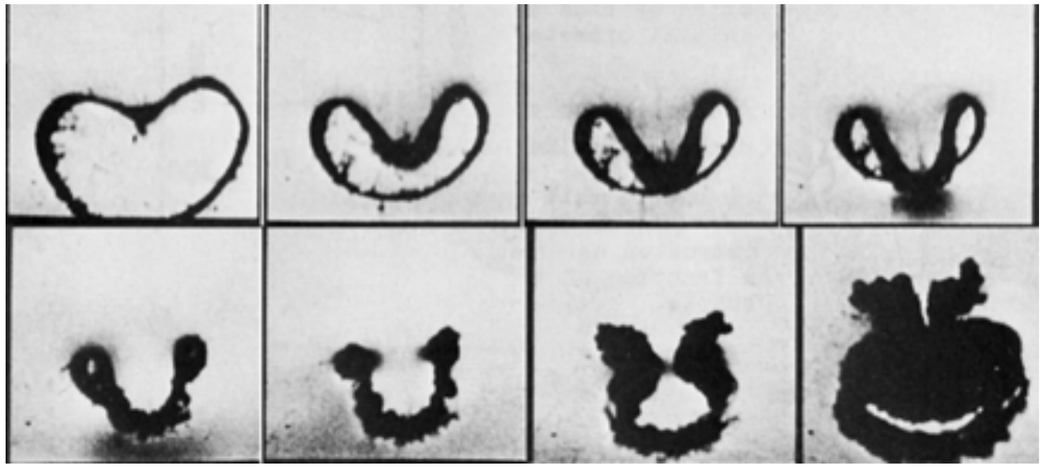

http://eswt.net/wp-content/uploads/2011/10/cavitation.gif

http://1.bp.blogspot.com/-uSYxEFf2_sw/Ukw-iDDivXI/AAAAAAAAIj8/X1Vm8thHLhA/s1600/5.+Stages+in+bubble+collapse.png

Topology of the cavities: genus (# of holes)

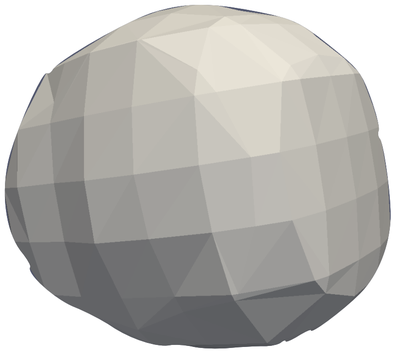

genus 0

genus 1

genus 2

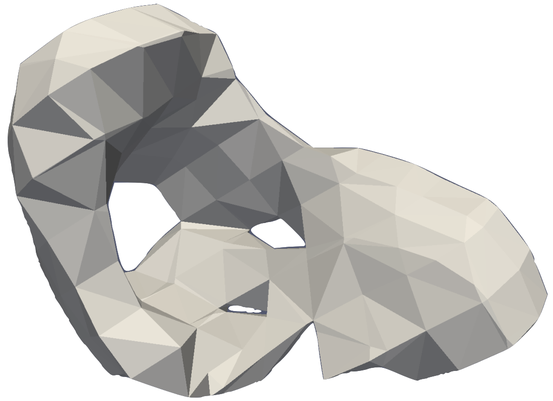

(genus 3)

caused by re-entrant micro-jet

Holes in the cavities from re-entrant micro-jets

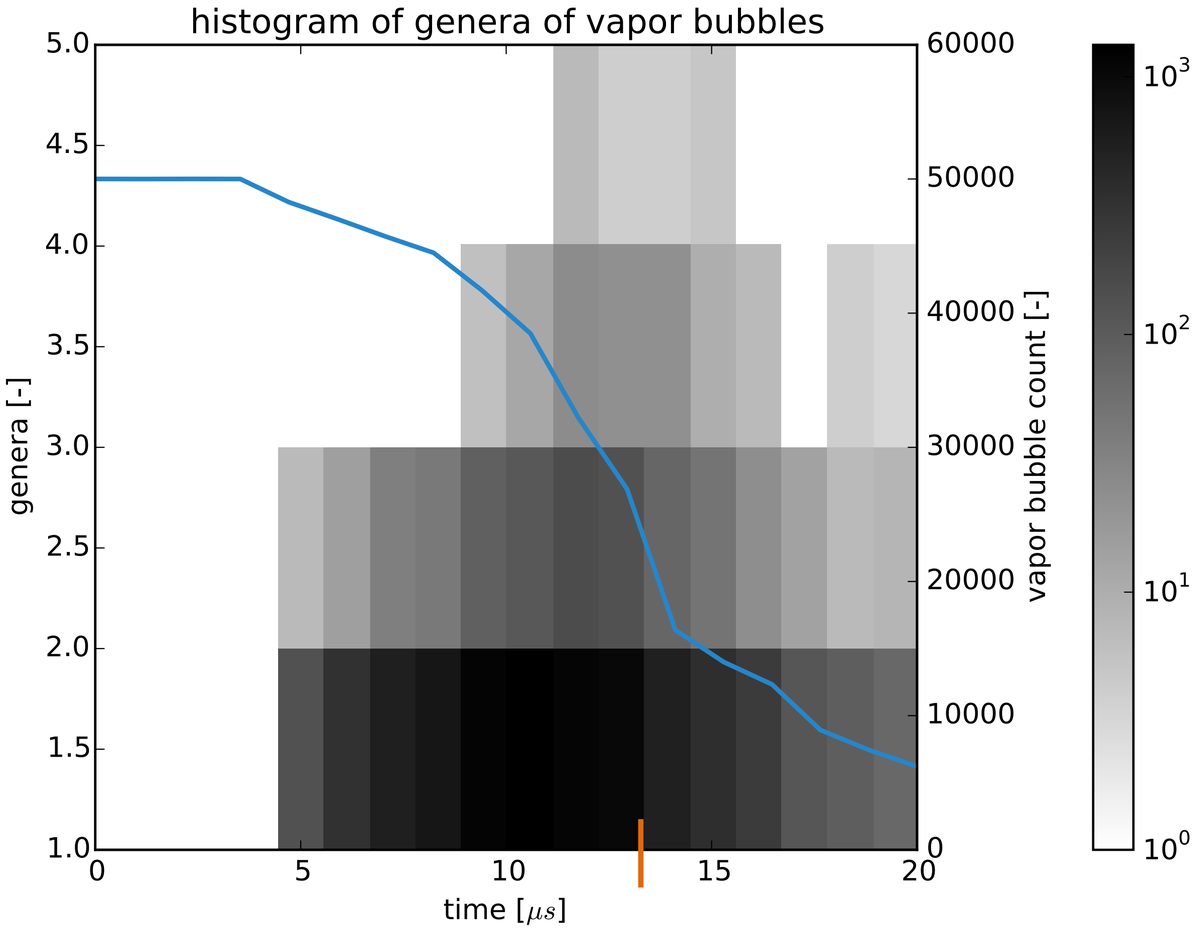

re-entrant micro-jets occur omnipresent before final collapse

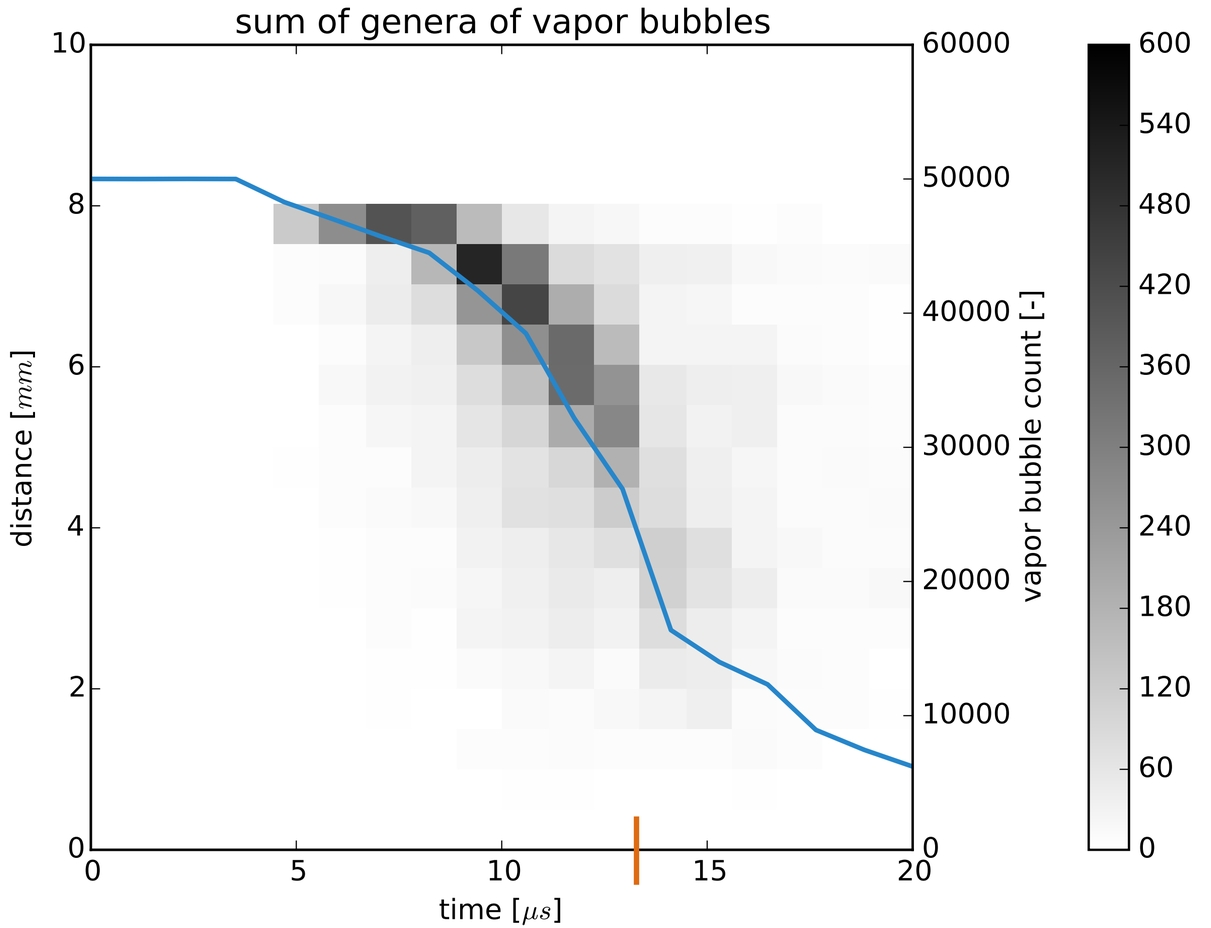

propagation of cavities with holes from cloud surface to center

ETH zürich

Cloud Cavitation Collapse

Fabian Wermelinger, Jonas Sukys, Christian Conti,
Panagiotis Hadjidoukas, Ursula Rasthofer, Diego Rossinelli,
Petros Koumoutsakos

Professorship for Computational Science, ETH Zurich, Switzerland

Uncertainty quantification
in cloud cavitation collapse

Collapse of two random clouds

2 clouds: different statistical realizations **(RNG seeds)** of the initial configuration

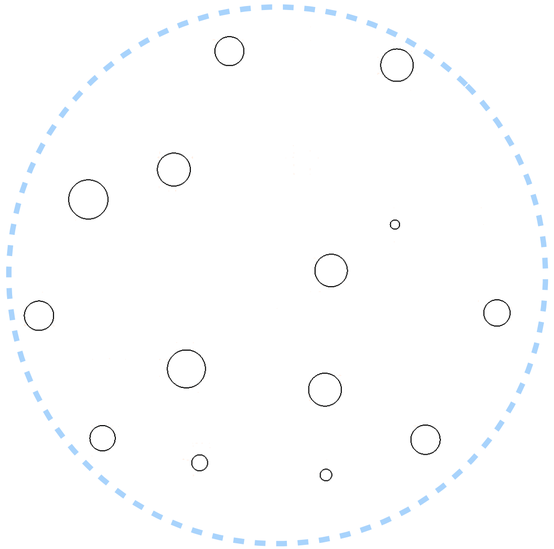

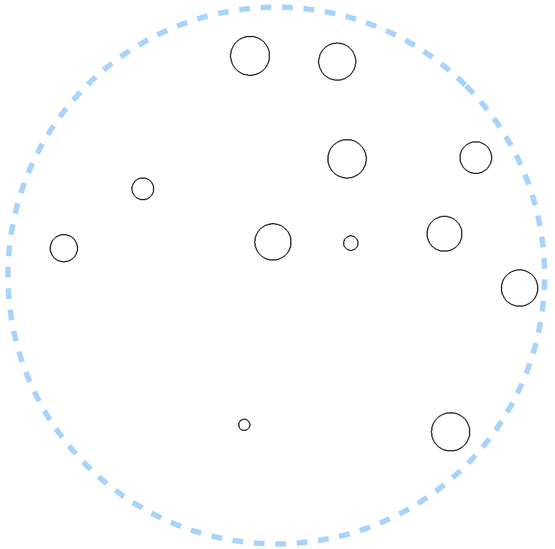

Spherical clouds of 100 equally sized (75μm) cavities

Uniformly distributed (random) cavity **positions**

Collapse of two random clouds

2 clouds: different statistical realizations **(RNG seeds)** of the initial configuration

Spherical clouds of 100 equally sized (75μm) cavities

Uniformly distributed (random) cavity **positions**

Multi-Level Monte Carlo [Heinrich, 1999] [Giles, 2008]

Variance reduction technique using sampling on a hierarchy of mesh resolutions

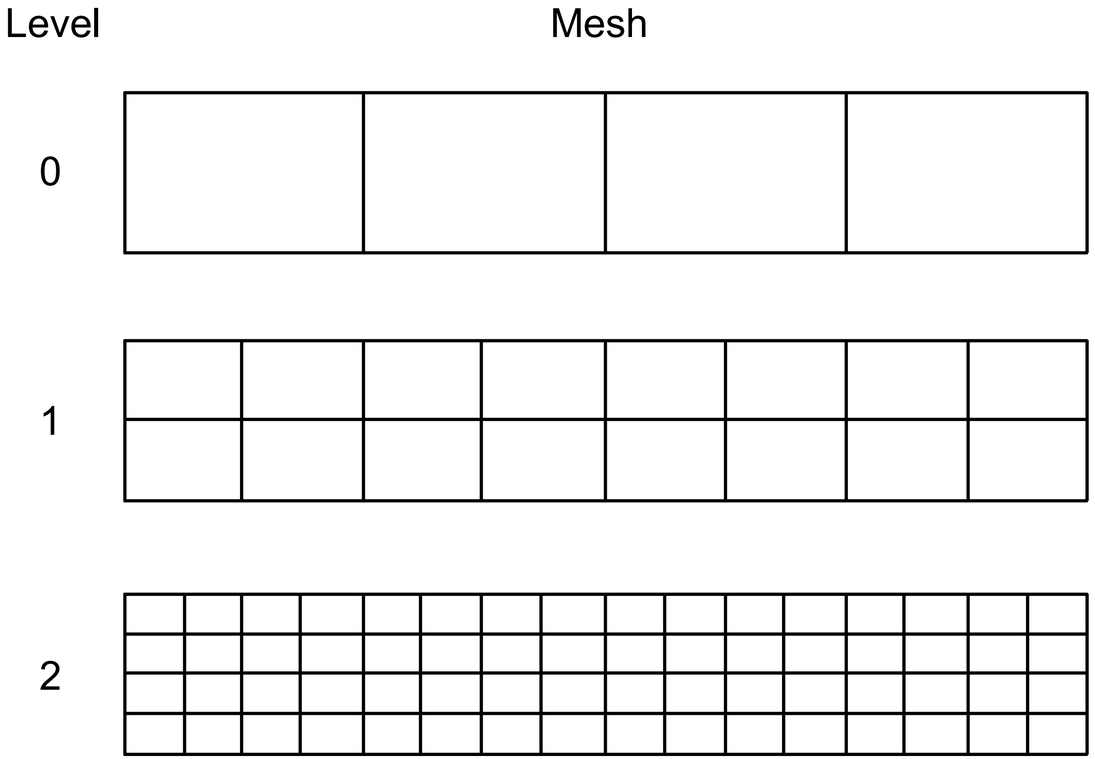

Multi-Level Monte Carlo method

Variance reduction technique using sampling on a hierarchy of mesh resolutions

1. Generate i.i.d. samples of random input quantities for each resolution level 0…L
2. For each level and sample, solve for approximate solutions using Cubism-MPCF
3. Assemble MLMC estimator for statistics of quantities of interest:

$$\mathbb{E}[q_L] \quad = \quad \mathbb{E}[q_0] \quad + \quad \sum_{\ell=1}^{L} \left(\mathbb{E}[q_\ell] - \mathbb{E}[q_{\ell-1}]\right) \quad \approx \quad \frac{1}{M_0}\sum_{i=1}^{M_0} q_0^i \quad + \quad \sum_{\ell=1}^{L} \frac{1}{M_\ell}\sum_{i=1}^{M_\ell} (q_\ell^i - q_{\ell-1}^i).$$

- Sampling **error** of the MLMC estimator is given in terms of **level correlations**:

$$\varepsilon^2 \quad = \quad \frac{\mathbb{V}[q_0]}{M_0} \quad + \quad \sum_{\ell=1}^{L} \frac{\mathbb{V}[q_\ell - q_{\ell-1}]}{M_\ell} \quad \approx \quad \mathbb{V}[q] \left(\frac{1}{M_0} \quad + \quad 2\sum_{\ell=1}^{L} \frac{1 - \mathrm{Cor}[q_\ell, q_{\ell-1}]}{M_\ell} \right).$$

Optimal control variate coefficients

- Each level in MLMC estimator is a special case of control variate with coefficient 1

$$\mathbb{E}[q_\ell] \quad \approx \quad \alpha\mathbb{E}[q_{\ell-1}] \quad + \quad \Big(\mathbb{E}[q_\ell] - \alpha\mathbb{E}[q_{\ell-1}]\Big).$$

- **Optimal** coefficient is given in terms of correlations between two levels

$$\alpha \quad = \quad \frac{\mathrm{Cov}[q_\ell, q_{\ell-1}]}{\mathrm{Var}[q_{\ell-1}]} \quad \approx \quad \mathrm{Cor}[q_\ell, q_{\ell-1}].$$

- This argument can be extended to the telescoping sum of the MLMC estimator

$$\mathbb{E}[q_L] \quad = \quad \alpha_0\mathbb{E}[q_0] \quad + \quad \sum_{\ell=1}^{L} \Big(\alpha_\ell\mathbb{E}[q_\ell] - \alpha_{\ell-1}\mathbb{E}[q_{\ell-1}]\Big).$$

- Related independent work for **reused** sampling on coarser levels:
  [Peherstorfer, Willcox, Gunzburger, 2015]

$$\alpha_\ell = \frac{\sigma^2_{L,\ell}}{\sigma^2_\ell}$$

Optimal control variate coefficients

Minimizes variance reduction costs for weakly correlated resolution levels

- Total **computational work-weighted** variance over **all** levels is given by

$$C[q_L^*] = \alpha_0^2 \mathbb{V}[q_0] W_0 + \sum_{\ell=1}^{L} \Big(\alpha_\ell^2 \mathbb{V}[q_\ell] + \alpha_{\ell-1}^2 \mathbb{V}[q_{\ell-1}] - 2\alpha_\ell \alpha_{\ell-1} \operatorname{Cov}[q_\ell, q_{\ell-1}] \Big) W_\ell .$$

- **Minimization** of the above pertains to solving **linear system** of equations,

$$\frac{\partial}{\partial \alpha_\ell} C[q_L^*] = 0, \qquad \ell = 0, \ldots, L-1.$$

- Linear system can be written in a form of a **diagonally dominant matrix**

$$\begin{bmatrix} \sigma_0^2 (W_1 + W_0) & -\sigma_{1,0}^2 W_1 & & & \\ -\sigma_{1,0}^2 W_1 & \ddots & \ddots & & \\ & \ddots & \ddots & -\sigma_{L-1,L-2}^2 W_{L-1} & \\ & & -\sigma_{L-1,L-2}^2 W_{L-1} & \sigma_{L-1}^2 (W_L + W_{L-1}) \end{bmatrix} \begin{bmatrix} \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_{L-2} \\ \alpha_{L-1} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ \sigma_{L,L-1}^2 W_L \end{bmatrix}$$

Example and comments

For two levels

- For two levels of resolution (i.e. L=1), optimal control variate coefficient is

$$\alpha_0 = \frac{W_1}{W_1 + W_0} \frac{\sigma_{1,0}}{\sigma_0}.$$

- Coarsest level already available — classical control variate coefficient is recovered
- Significantly more expensive and **strongly** correlated finer level — classical MLMC
- **WARNING:** significantly more expensive but **weakly** correlated finer level with

$$\rho_{1,0} < \frac{1}{2} \frac{W_1 + W_0}{W_1}.$$

leads to **variance increase** in MLMC, unless optimal control variates are used

Optimized number of samples [Giles, 2008]

Using empirical estimators for variances and measurements of computations work

- Sampling **error** of the MLMC estimator is given in terms of **level variances**:

$$\varepsilon^2 = \frac{\mathbb{V}[\alpha_0 q_0]}{M_0} + \sum_{\ell=1}^{L} \frac{\mathbb{V}[\alpha_\ell q_\ell - \alpha_{\ell-1} q_{\ell-1}]}{M_\ell} \approx \frac{\tilde{\sigma}_0^2}{M_0} + \sum_{\ell=1}^{L} \frac{\tilde{\sigma}_\ell^2}{M_\ell}.$$

Optimization problem

Given a required tolerance τ and variances σ_ℓ^2 each level, minimize computational work and find optimal number of samples such that tolerance is attained: $\varepsilon \leq \tau$.

Optimized number of samples

Using Lagrange multipliers for derivations, optimized number of samples are given by

$$M_\ell = \left\lceil \frac{1}{\tau^2} \sqrt{\frac{\tilde{\sigma}_\ell^2}{W_\ell}} \sum_{k=0}^{L} \sqrt{\tilde{\sigma}_k^2 W_k} \right\rceil .$$

Remark: an analogous result is available for a prescribed computational budget (instead of tolerance).

PyMLMC

GIT repository: pymlmc.sukys.lt

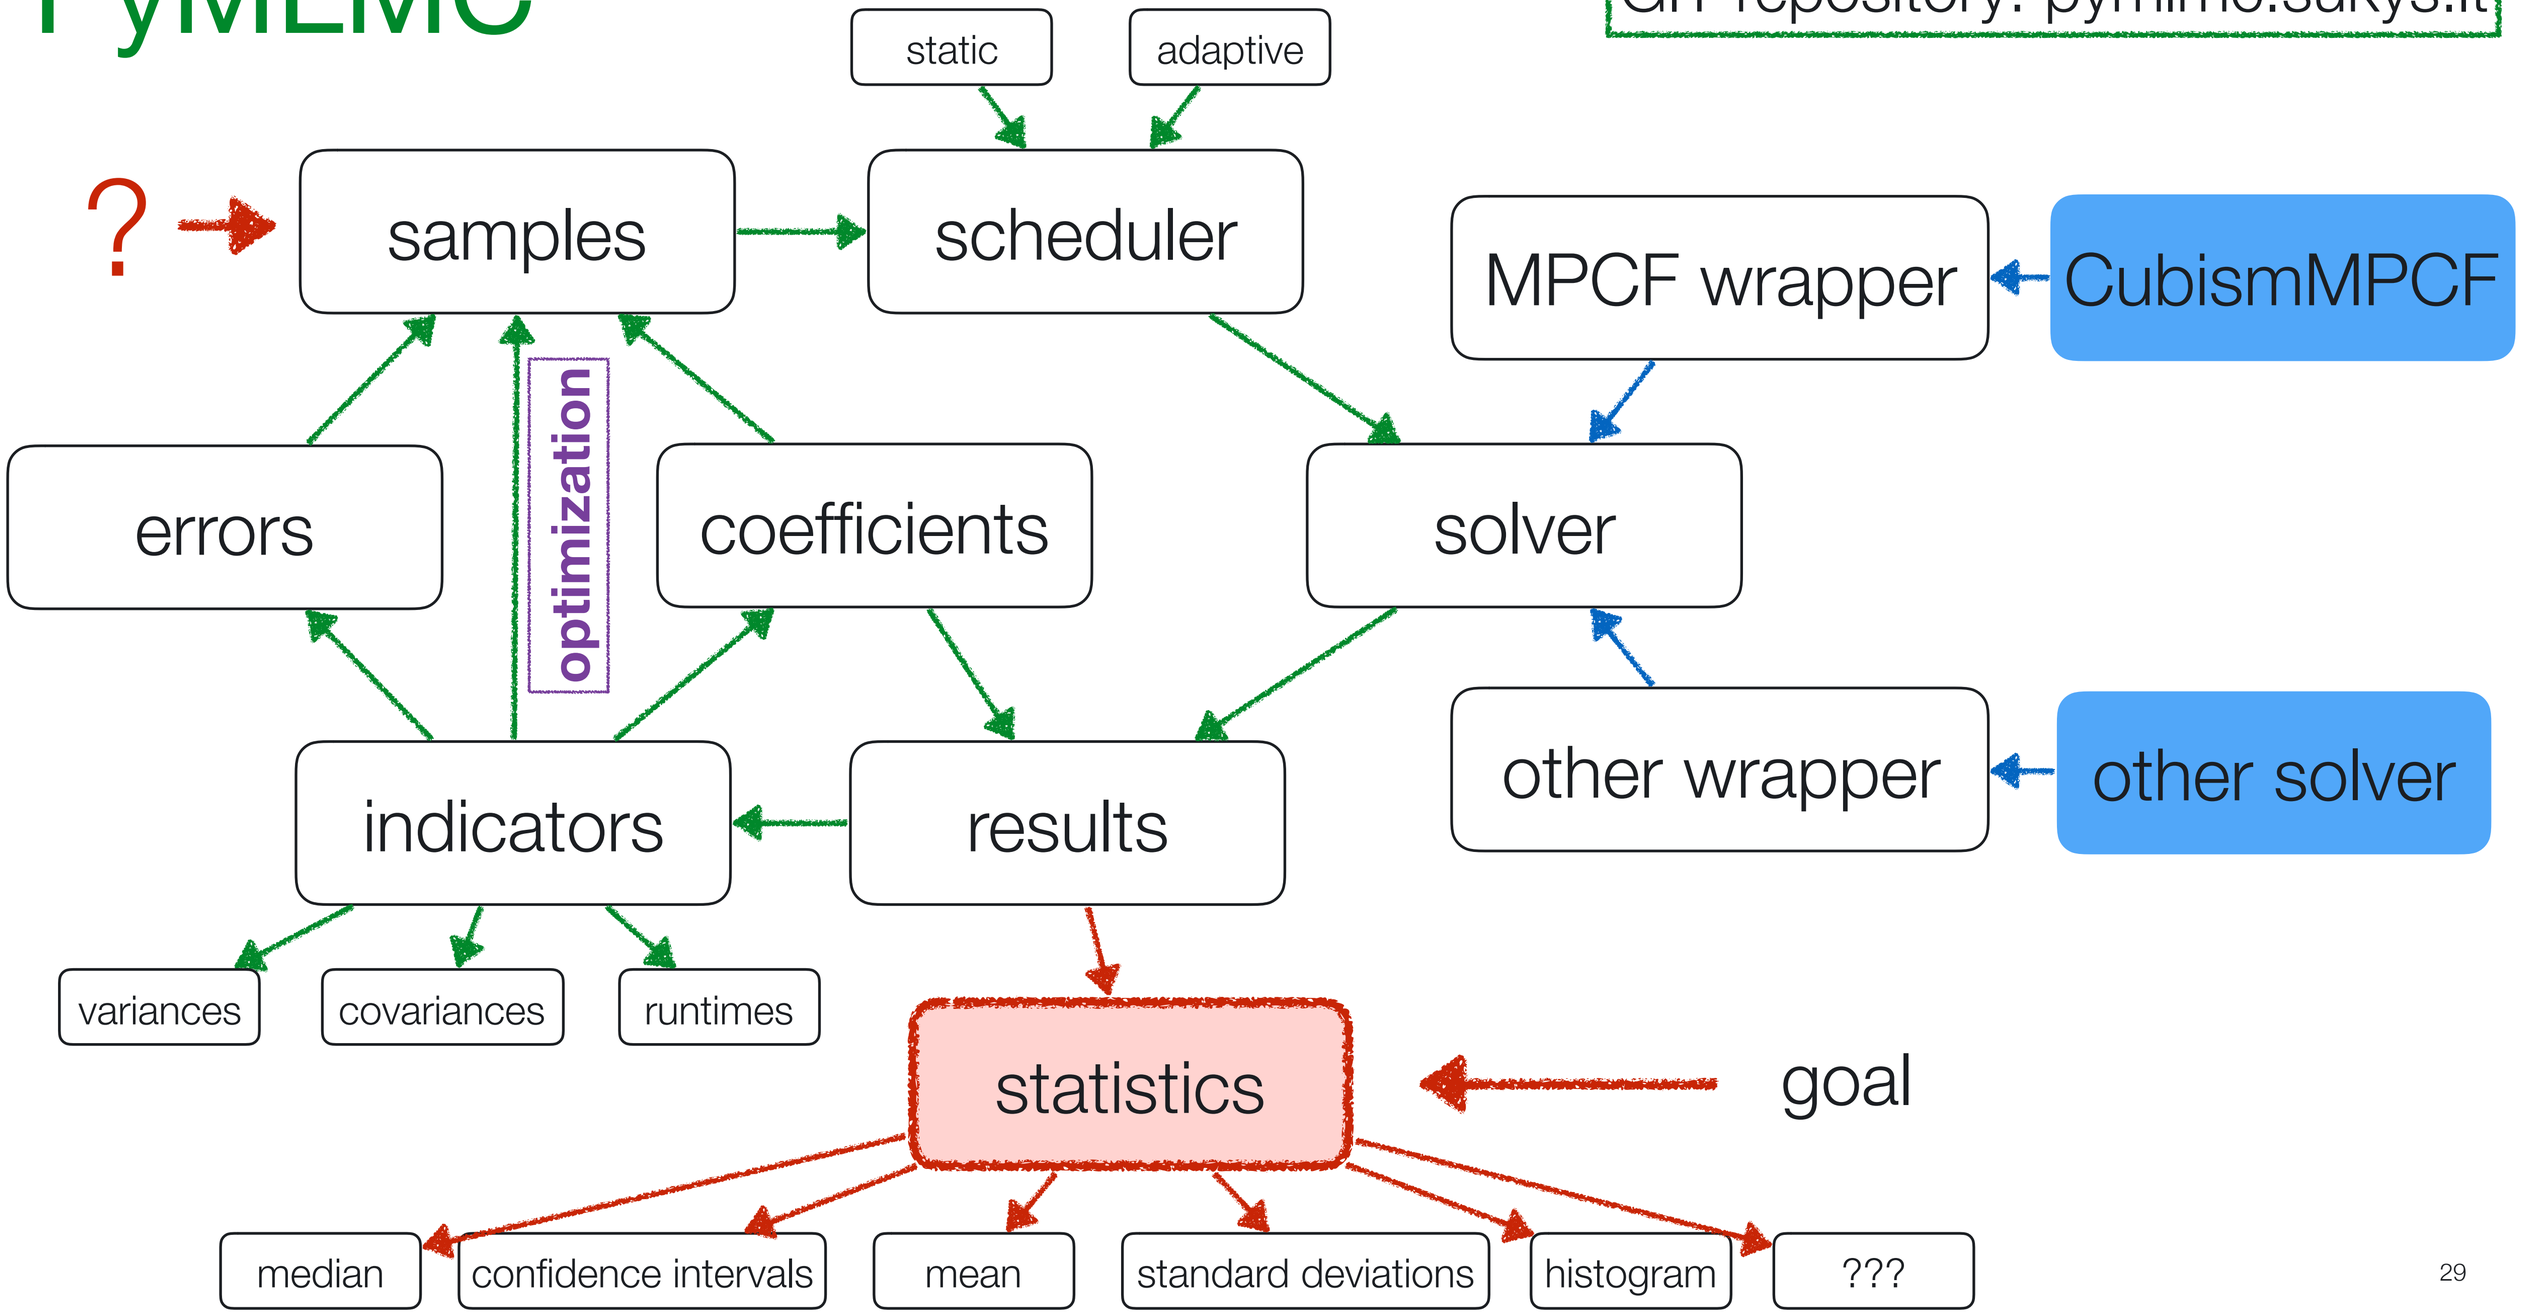

Insight to inner workings of MLMC

Majority of samples computed on lowest levels of resolution - reduced budget

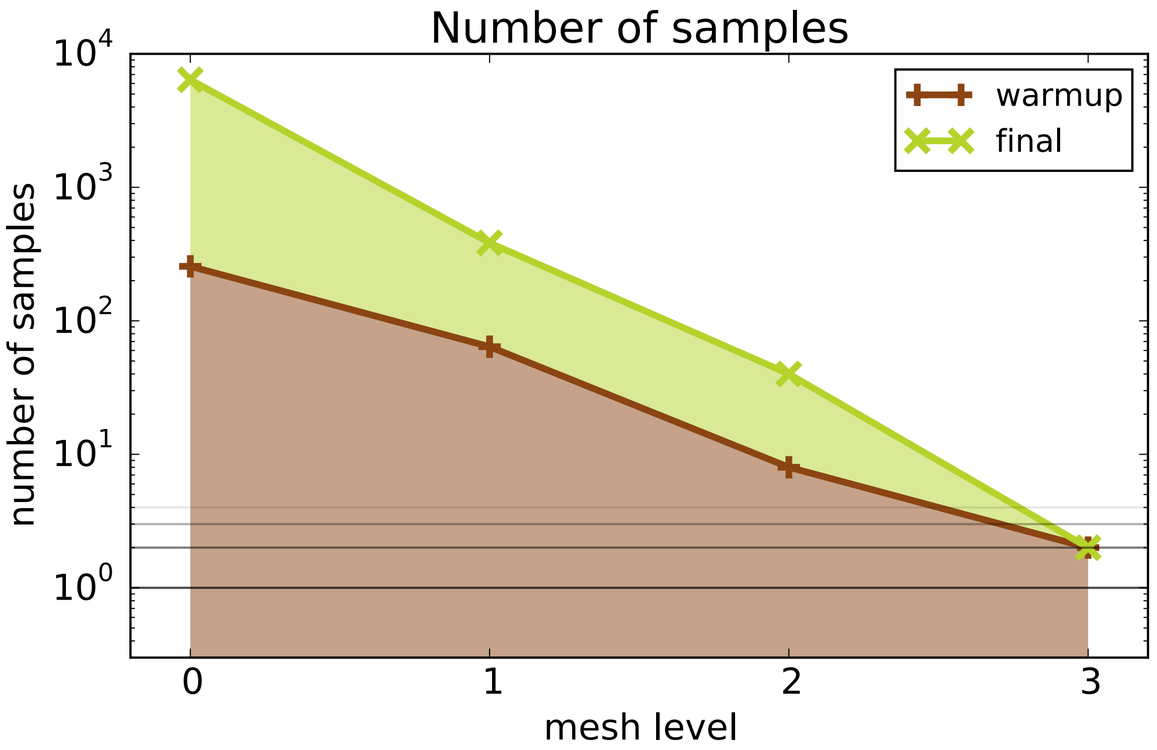

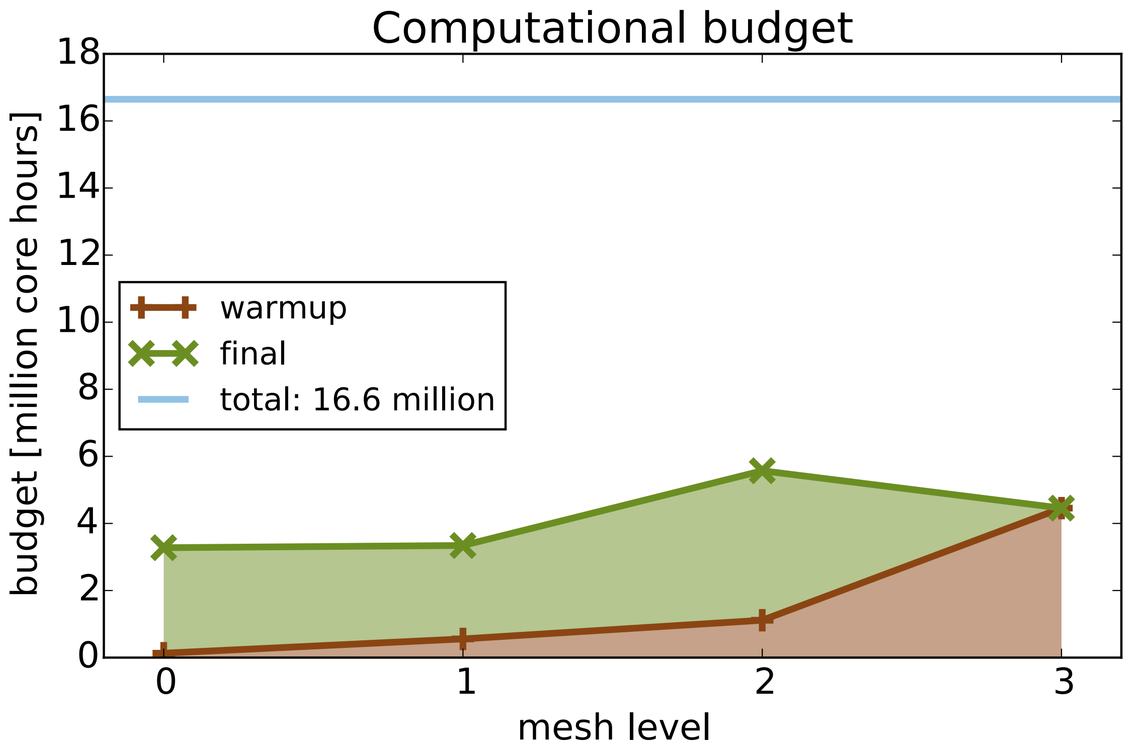

adaptive number of warmup samples

observed speedup: 176x

Insight to inner workings of MLMC

Optimal control variate coefficients

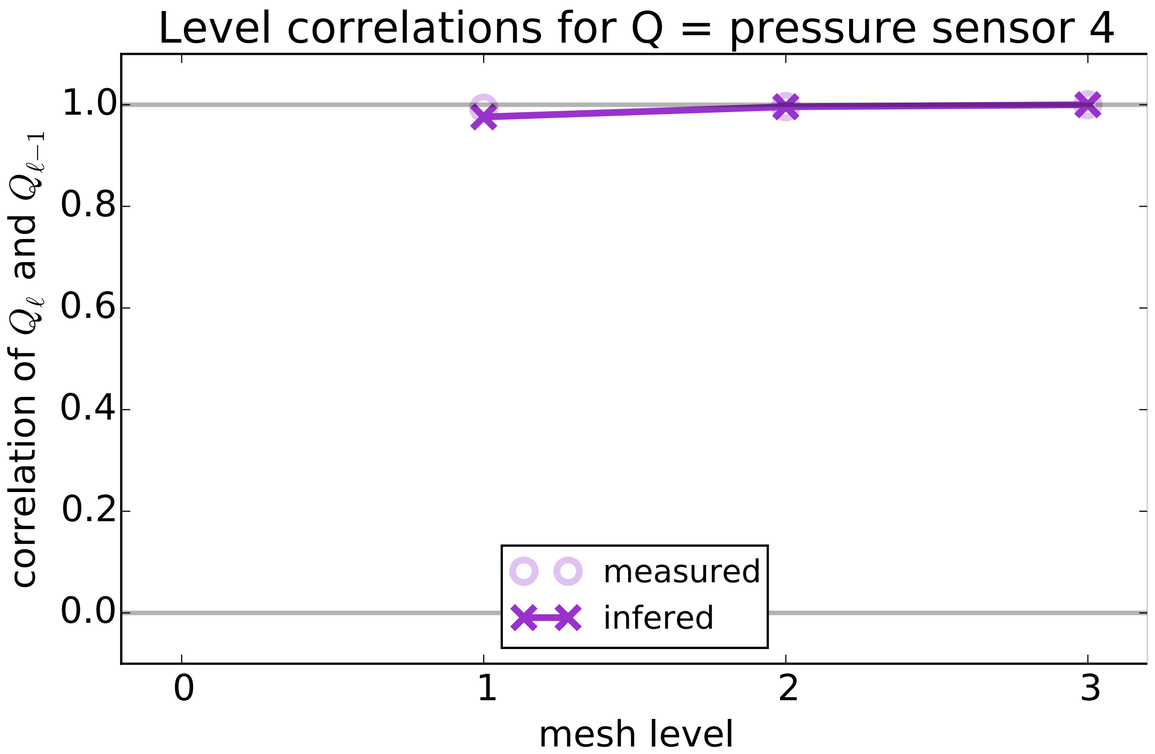

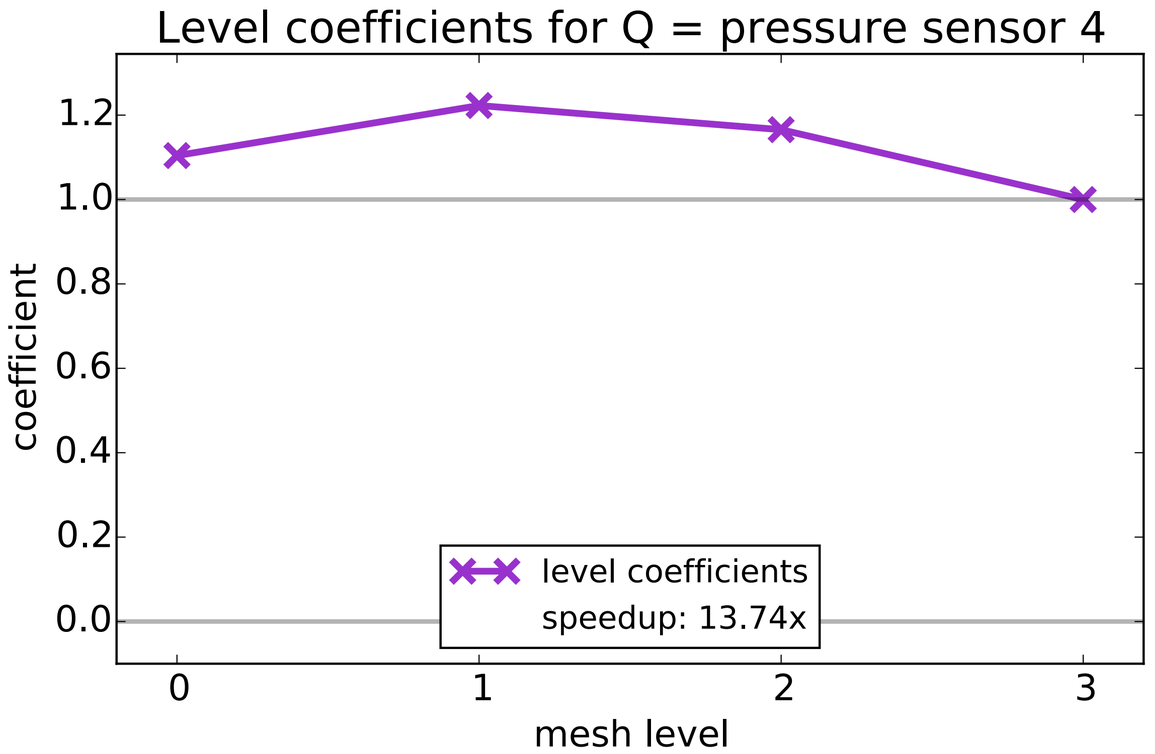

Insight to inner workings of MLMC

Correlation estimates for differences between resolution levels decrease

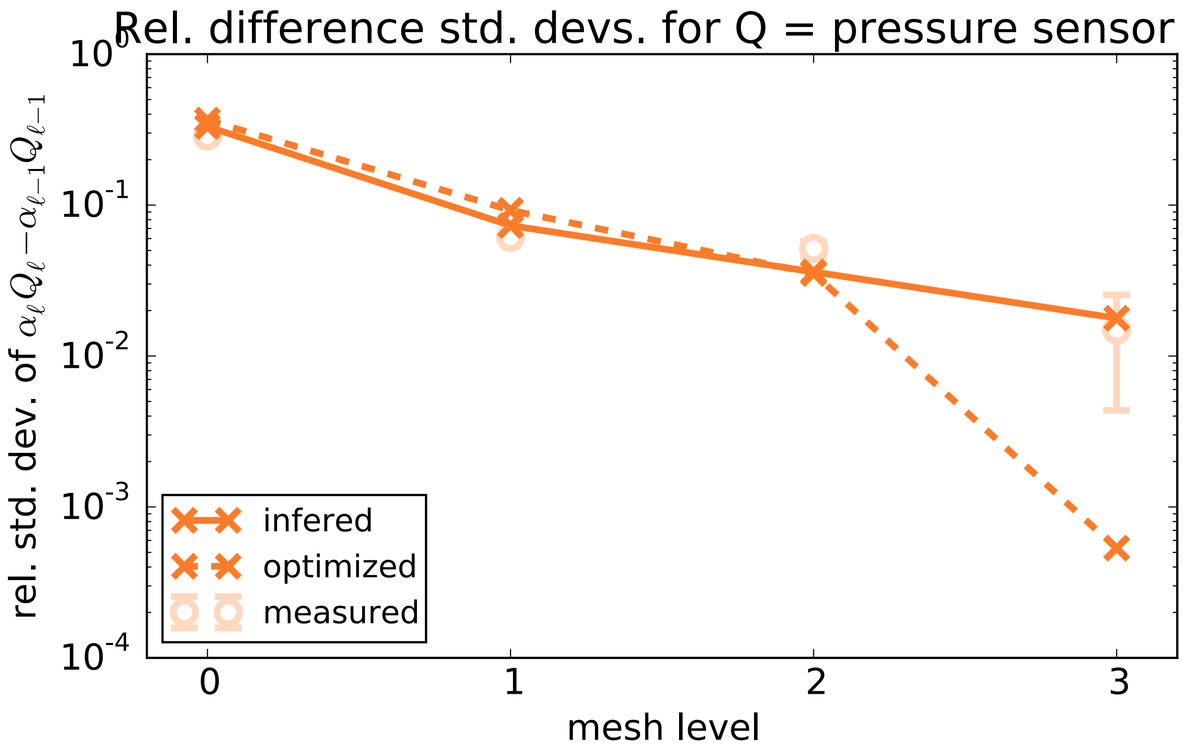

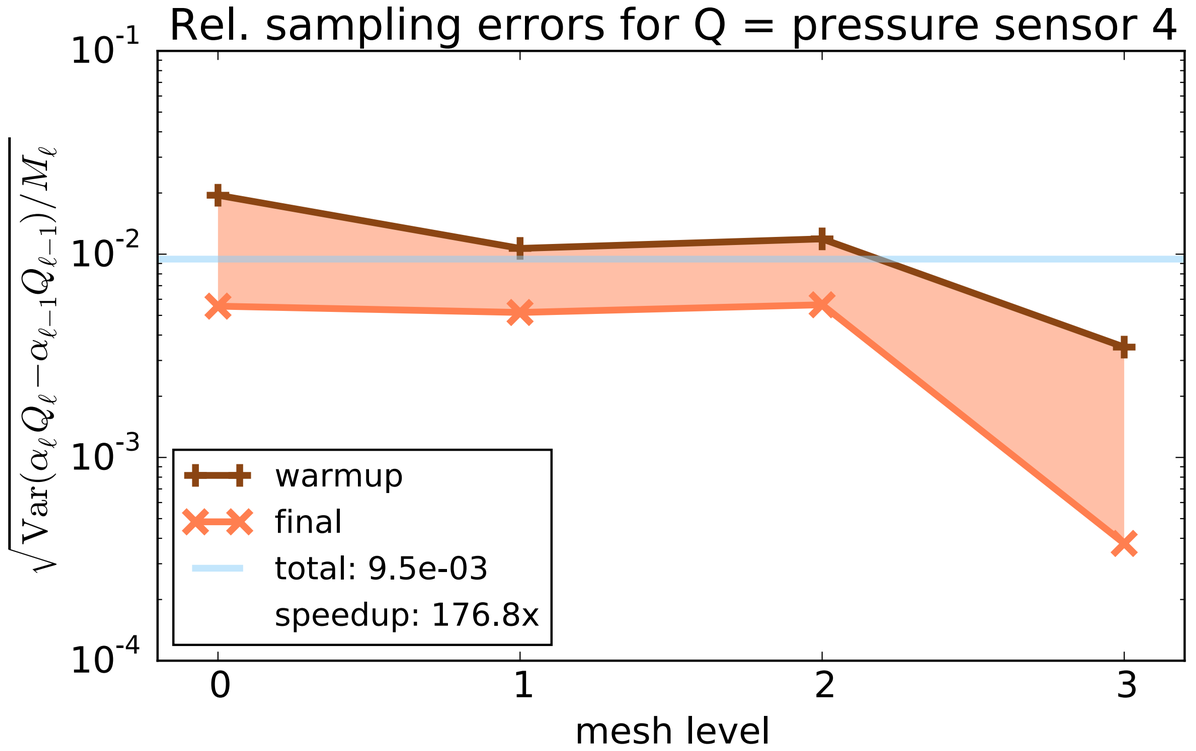

Results of MLMC

Uncertainty quantification (i.e. mean, confidence intervals) for QoIs

vapor volume

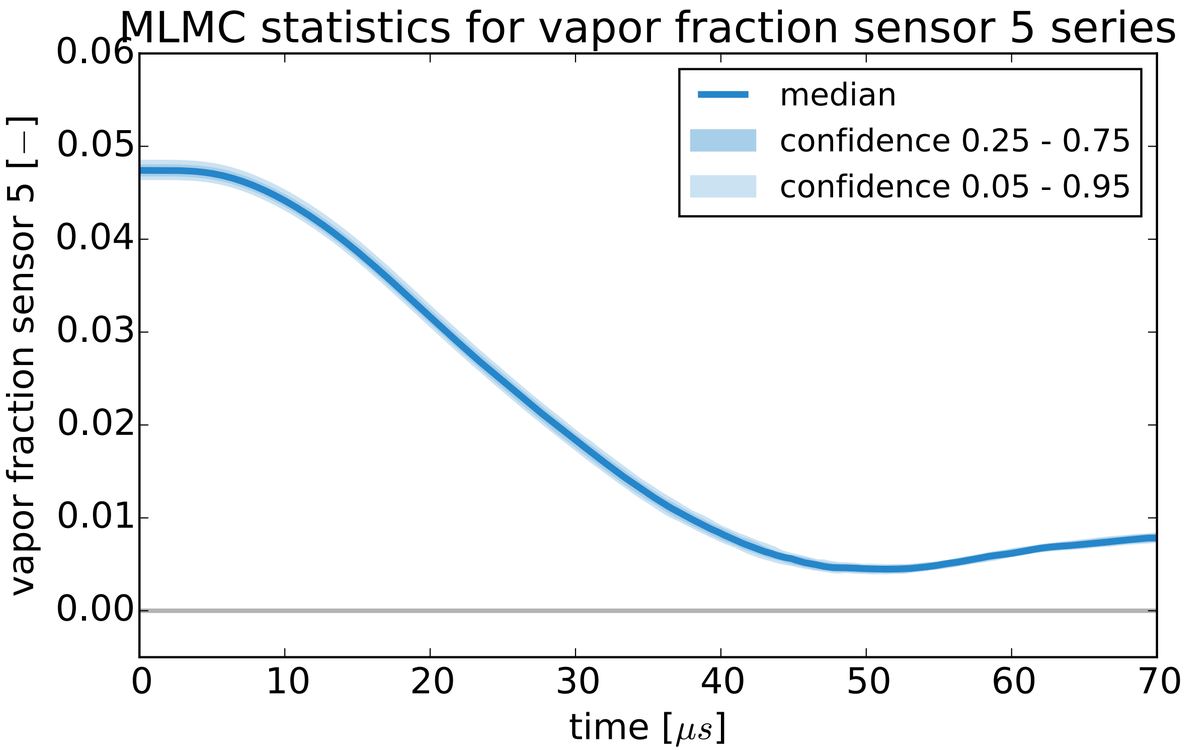

no significant uncertainty

pressure sensor

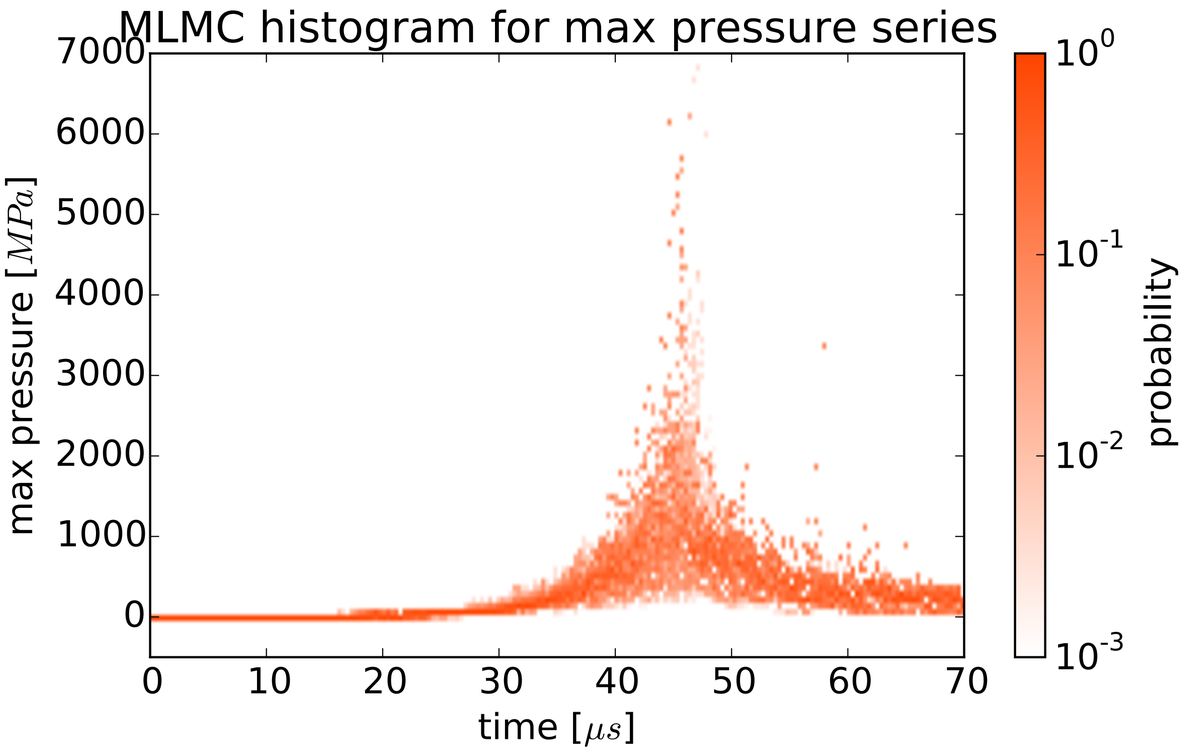

wide range of probable pressures

100 MPa - 6000 MPa

Results of MLMC

Evolution of inwards propagating pressure wave

peak pressure location within the cloud

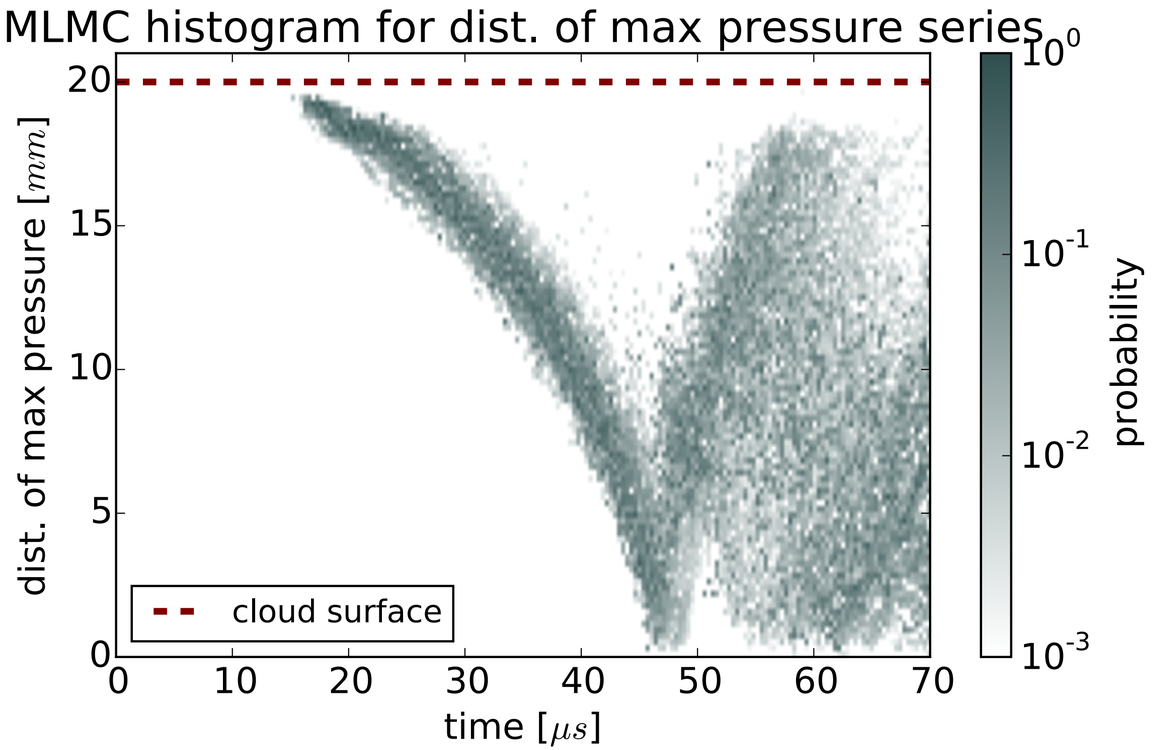

Results of MLMC

Secondary cavitation observed at the epicenter immediately after the final collapse

vapor volume fraction sensor

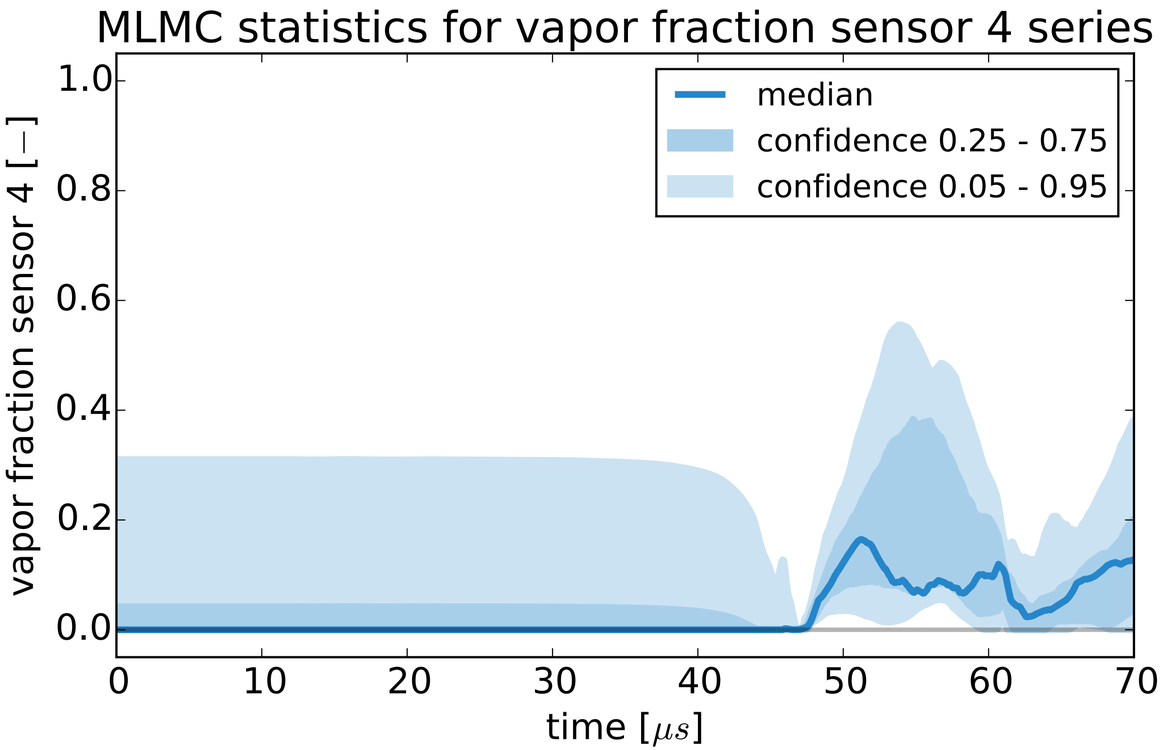

“secondary cavitation” region
after the final cloud collapse

pressure sensor

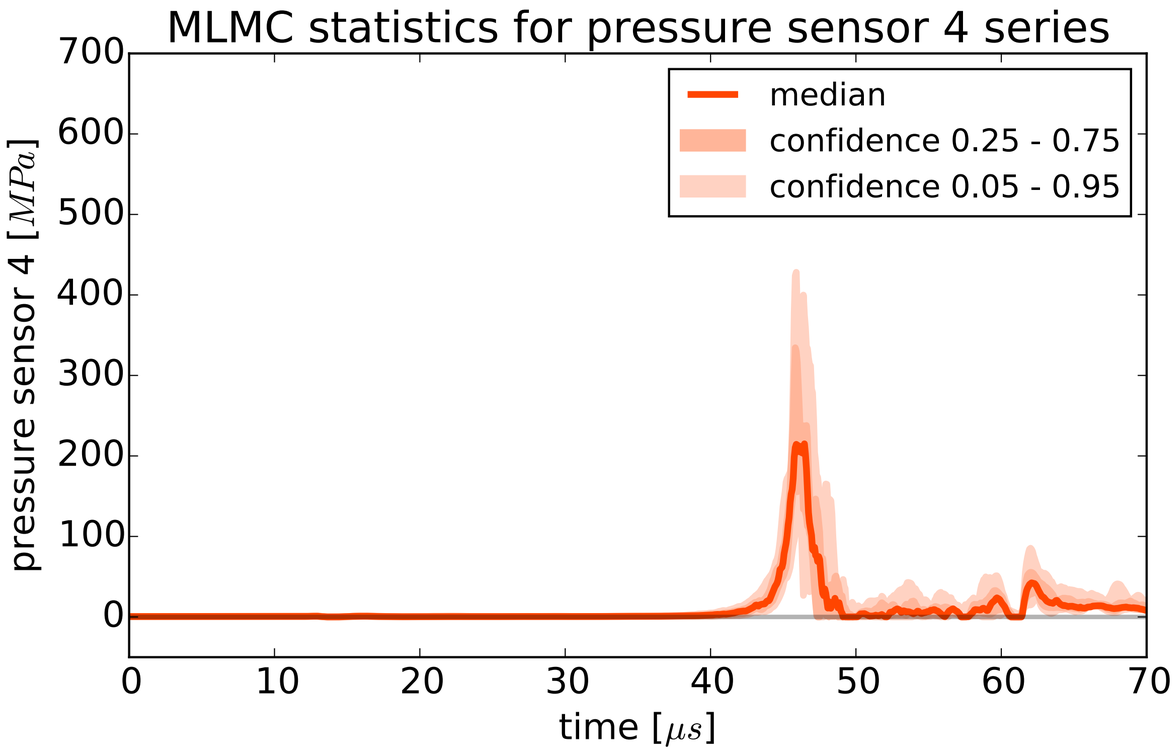

wide 90% confidence interval
100 MPa - 500 MPa

Two dimensional slices

Single realization of a slice through the center of the cloud

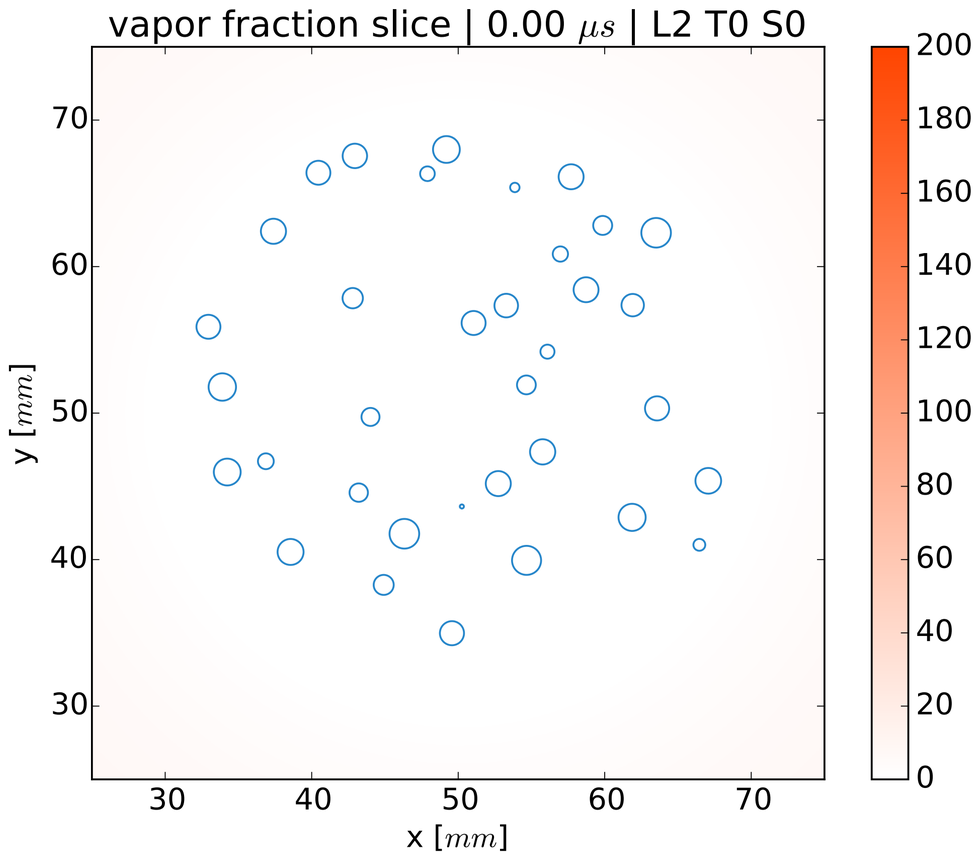

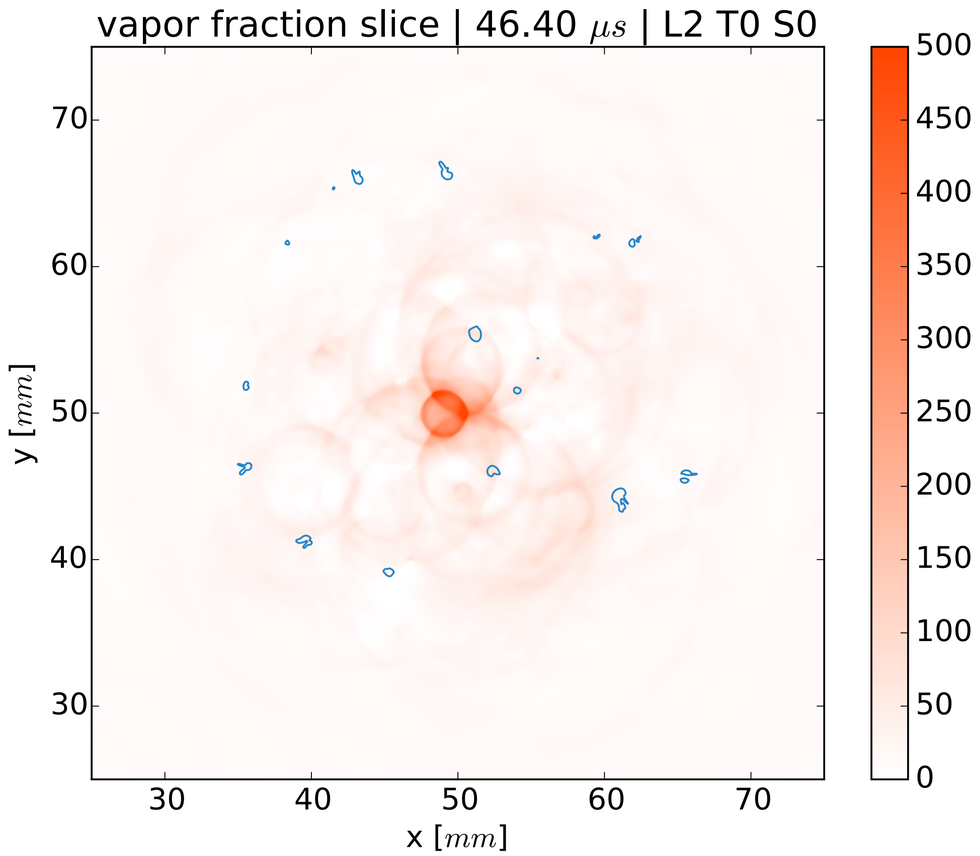

One dimensional lines

Uncertainties along the line through the center of the cloud **(at peak collapse)**

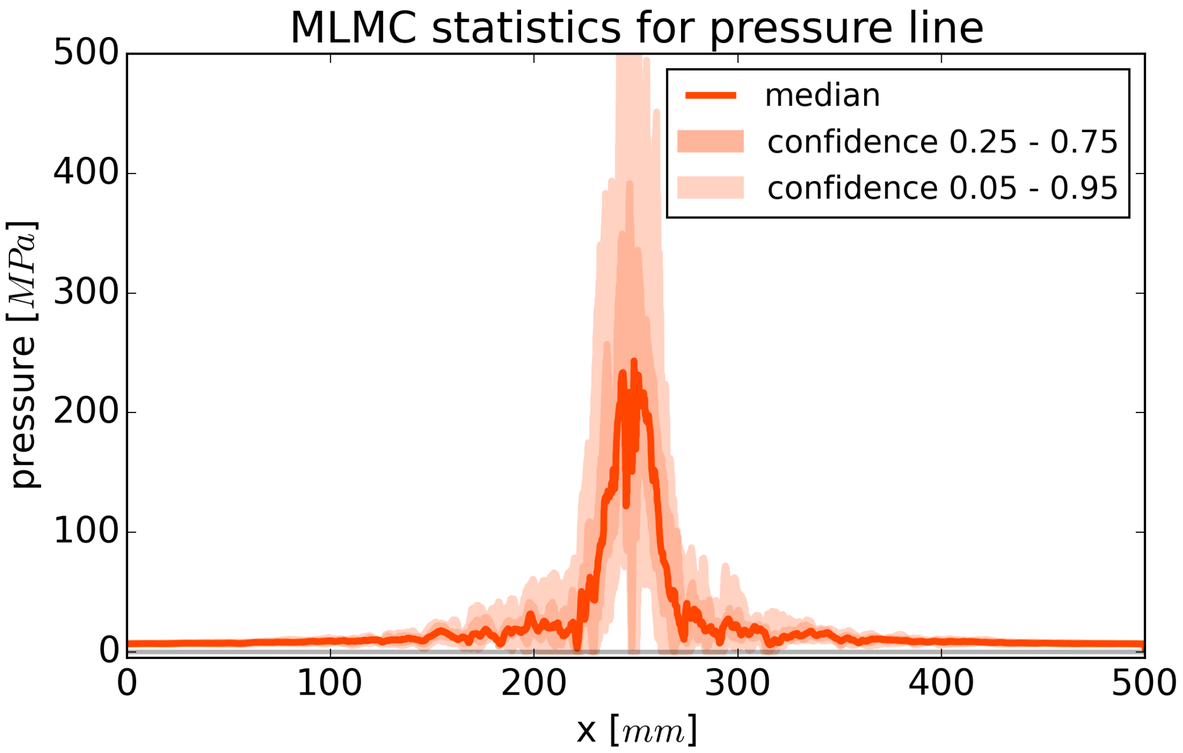

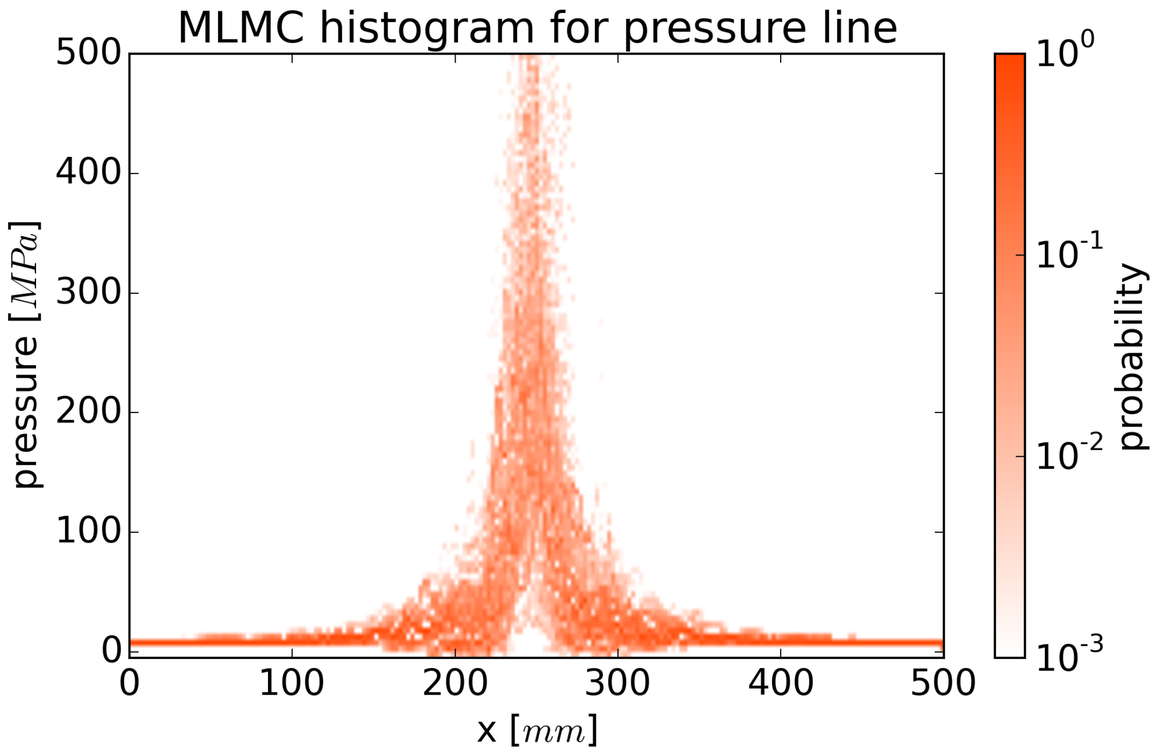

Optimal control variate coefficients

Speedup

	Level samples	Budget in CPU hours	Total speedup	Relative error
MC	∞	2 billion	-	9.5E-03
MLMC	4352, 258, 32, 3	50 million	50.6	9.5E-03
OCV-MLMC	6400, 384, 40, 2	16.6 million	176.8	9.5E-03

higher OCV-MLMC speedup is expected
for even less correlated levels

Summary and outlook

- **OCV-MLMC** for uncertainty quantification in multiple **sensors for pressure**, density, etc.
  - instead of a single value, **confidence intervals or PDFs for peak pressures** are provided
- **Optimal control variate coefficients** for weakly correlated levels without sample “recycling”
- **Fault tolerance:** if some samples fail, the rest are used to assemble estimators [Pauli, Schwab, Arbenz]
- **OUTLOOK**
  - improve quality of constructed PDFs, especially on finer levels [collaboration with T. Barth, NASA]
  - discrete optimization for the number of samples [Pauli, Arbenz]
  - unbiased estimator using randomised resolution levels [Rhee, Glynn, 2015]
  - investigate the causes for large uncertainties in the peak pressures during cloud cavitation collapse

Improving the quality of PDFs using KDEs

Optimal kernel widths for kernel density estimators of the PDFs

few samples!

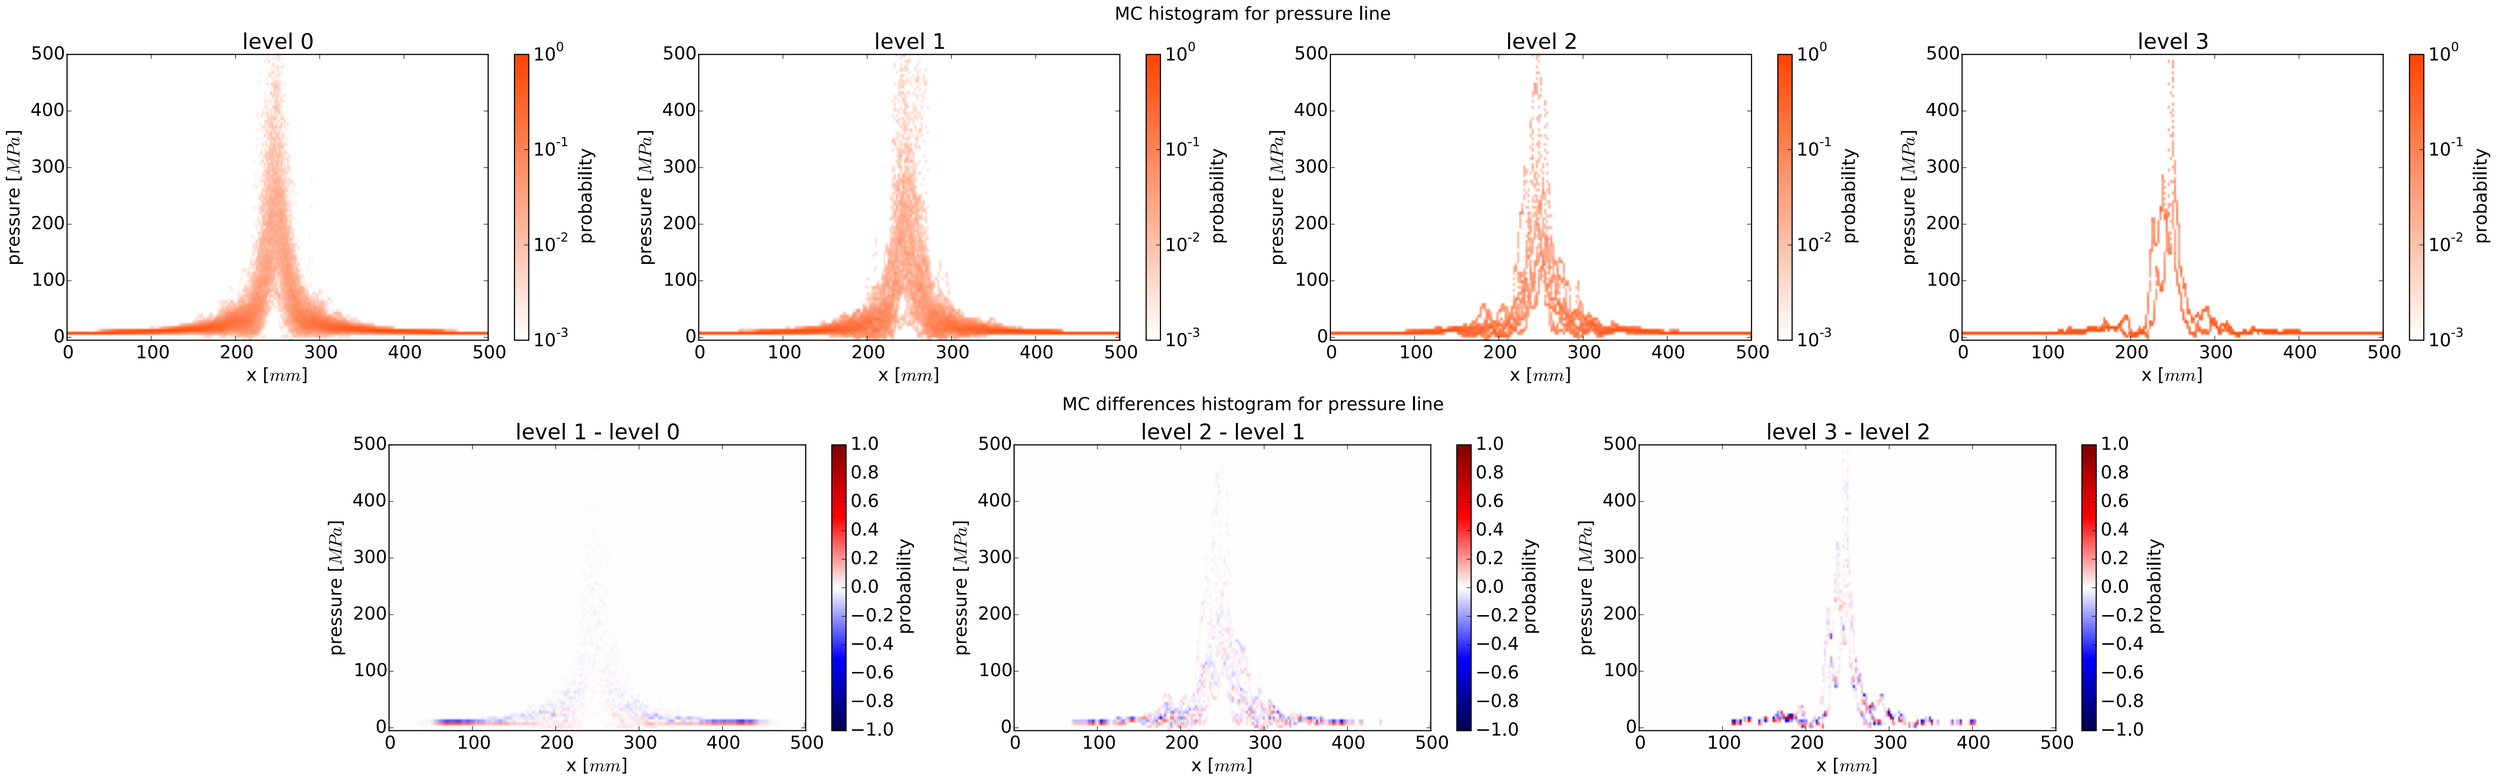

Idea: propagate information about optimal kernel widths from coarser levels to finer levels by means of **Bayesian inference**

HPC resources

CSCS allocation
Project s500
Piz Daint
Cray XC30
42 176 cores
5 272 GPUs
7.8 PFlops
Switzerland

INCITE allocation
Argonne National Labs
Project “CloudPredict”
MIRA
BlueGene/Q
786 432 cores
10 PFlops
United States

PRACE allocation
Jülich Research Center
Project 091
JUQUEEN
BlueGene/Q
458 752 cores
5.9 PFlops
Germany

PRACE allocation
CINECA
Project 09_2376
FERMI
BlueGene/Q
163 840 cores
2.1 PFlops
Italy

Thanks to


ETH zürich


CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre


Argonne
NATIONAL LABORATORY


FNSNF
Fonds national suisse
Schweizerischer Nationalfonds
Fondo nazionale svizzero
Swiss National Science Foundation


U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING


CINECA


IBM


PRACE


JÜLICH
FORSCHUNGSZENTRUM

Selected publications

Most recent list available in: pub.sukys.lt

- J. Šukys, U. Rasthofer, F. Wermelinger, P. Hadjidoukas, D. Rossinelli, P. Koumoutsakos. Numerical investigation of cavitation dynamics of O(10^4) bubble clouds: cloud geometries, identification of characteristic stages, and cavity classification. In progress, 2016.
- S. Mishra, Ch. Schwab and J. Šukys. Multi-level Monte Carlo Finite Volume methods for uncertainty quantification of acoustic wave propagation in random heterogeneous layered medium. **J. Comput. Phys.**, 312:192-217, 2016.
- C.S. Linares, M.J. Castro, S. Mishra and J. Šukys. MLMC finite volume method for shallow water equations with uncertain parameters applied to landslide-generated tsunamis. **Applied Mathematical Modelling**, 39:7211-7226, 2015.
- J. Šukys. Adaptive load balancing for massively parallel multi-level Monte Carlo solvers. Parallel Processing and Applied Mathematics 2013, Part I, **Springer Lecture Notes in Computational Science** 8384:47-56, 2014. Springer, Berlin Heidelberg.
- S. Mishra, Ch. Schwab, and J. Šukys. Multi-level Monte Carlo finite volume methods for nonlinear systems of conservation laws in multi-dimensions. **Journal of Computational Physics**, 231(8):3365-3388, 2012.
- S. Mishra, Ch. Schwab, and J. Šukys. Multi-level Monte Carlo Finite Volume methods for shallow water equations with uncertain topography in multi-dimensions. **SIAM Journal of Scientific Computing**, 34(6):B761-B784, 2012.
- S. Mishra, Ch. Schwab, and J. Šukys. Multi-level Monte Carlo finite volume methods for uncertainty quantification in nonlinear systems of balance laws. **Lecture Notes in Computational Science and Engineering** 92, 2013.
- P. Hadjidoukas, D. Rossinelli, F. Wermelinger, J. Šukys, U. Rasthofer, C. Conti, Hejazialhosseini, Koumoutsakos. High throughput simulations of two-phase flows on BlueGene/Q. **Advances in Parallel Computing**, IOS Press, 2016.
- A. Barth, Ch. Schwab and J. Šukys. Multi-level Monte Carlo approximations of statistical solutions to the Navier-Stokes equation. MCQMC 2014, **Springer Proceedings in Mathematics & Statistics** 163, 2016.
- J. Šukys, Ch. Schwab, and S. Mishra. Multi-level Monte Carlo Finite Difference and Finite Volume methods for stochastic linear hyperbolic systms. MCQMC 2012, **Springer Proceedings in Mathematics & Statistics**, 65:649-666, 2013.
- J. Šukys, S. Mishra, and Ch. Schwab. Static load balancing for multi-level Monte Carlo finite volume solvers. Parallel Processing and Applied Mathematics 2011, Part I, **Springer Lecture Notes in Computational Science** 7203:245-254, 2012. Springer, Heidelberg.
- …

THANK YOU

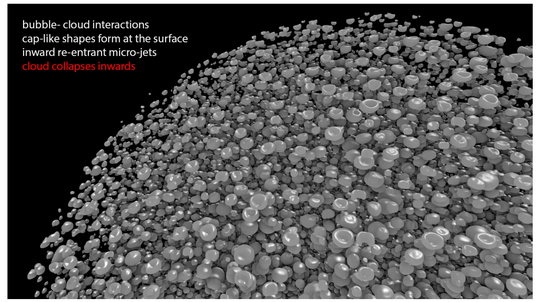
bubble- cloud interactions
cap-like shapes form at the surface
inward re-entrant micro-jets
cloud collapses inwards

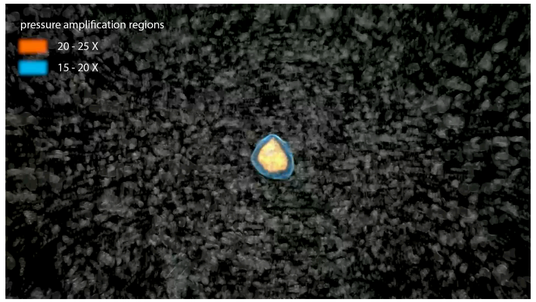
pressure amplification regions
20 - 25 X
15 - 20 X